

Bayesian CART models for insurance claims frequency

Yaojun Zhang, Lanpeng Ji^{*}, Georgios Aivaliotis, Charles Taylor

Department of Statistics, University of Leeds, United Kingdom

ARTICLE INFO

JEL classification:

C11
C14
C51
G22

Keywords:

Bayesian CART
Claims frequency
DIC
Insurance pricing
MCMC
Negative binomial distribution
Zero-inflated Poisson distribution

ABSTRACT

The accuracy and interpretability of a (non-life) insurance pricing model are essential qualities to ensure fair and transparent premiums for policy-holders, that reflect their risk. In recent years, classification and regression trees (CARTs) and their ensembles have gained popularity in the actuarial literature, since they offer good prediction performance and are relatively easy to interpret. In this paper, we introduce Bayesian CART models for insurance pricing, with a particular focus on claims frequency modelling. In addition to the common Poisson and negative binomial (NB) distributions used for claims frequency, we implement Bayesian CART for the zero-inflated Poisson (ZIP) distribution to address the difficulty arising from the imbalanced insurance claims data. To this end, we introduce a general MCMC algorithm using data augmentation methods for posterior tree exploration. We also introduce the deviance information criterion (DIC) for tree model selection. The proposed models are able to identify trees which can better classify the policy-holders into risk groups. Simulations and real insurance data will be used to illustrate the applicability of these models.

1. Introduction

An insurance policy refers to an agreement between an insurance company (the insurer) and a policy-holder (the insured), in which the insurer promises to charge the insured a certain fee for some unpredictable losses of the customer within a period of time, usually one year. The charged fee is called a *premium* which includes a *pure premium* and other loadings such as operational costs. For each policy, the pure premium is determined by multiple explanatory variables (such as characteristics of the policy-holders, the insured objects, the geographical region, etc.), also called *risk factors* (Ohlsson and Johansson, 2010). The premium charged reflects the customer's degree of risk; a higher premium suggests a potential higher risk, and vice versa. Therefore, it is necessary to use risk factors to classify policy-holders with similar risk profiles into the same tariff class. The insureds in the same group, all having similar risk characteristics, will pay the same reasonable premium. The process of constructing these tariff classes is also known as *risk classification*; see, e.g., Denuit et al. (2007); Henckaerts et al. (2018). In the basic formula of non-life insurance pricing, the pure premium is obtained by multiplying the expected claims frequency with the conditional expectation of severity, assuming independence between frequency and severity; see, e.g., Henckaerts et al. (2021). Hence, modelling the claims frequency represents an essential first step in non-life insurance pricing. In this paper, we propose efficacious approaches (namely, Bayesian CARTs or BCART models) to analyze imbalanced insurance claims frequency data.

Due to its flexibility in modelling a large number of distributions in the exponential family, generalized linear models (GLMs), developed in Nelder and Wedderburn (1972), have been the industry-standard predictive models for insurance pricing (Denuit et al., 2007; Wuthrich, 2022). Explanatory variables enter a GLM through a linear predictor, leading to interpretable effects of the risk factors on the response. Extensions of GLMs to generalized additive models (GAMs) to capture the nonlinear effects of risk factors sometimes offer more flexible models. However, both GLMs and GAMs often fail to identify the complex interactions among risk factors. Another popular classical method based on Bayesian statistics, the credibility method, was introduced to deal with multi-level factors and lack of data issues; see, e.g., Ohlsson and Johansson (2010); Bühlmann and Gisler (2005). Because of the limitations of these classical statistical methods and equipped with continually developing technologies, further research has recently turned to machine learning techniques. Several machine learning methods such as neural networks, regression trees, bagging

^{*} Corresponding author.

E-mail addresses: mmyz@leeds.ac.uk (Y. Zhang), lji@leeds.ac.uk (L. Ji), G.Aivaliotis@leeds.ac.uk (G. Aivaliotis), c.c.taylor@leeds.ac.uk (C. Taylor).

<https://doi.org/10.1016/j.insmatheco.2023.11.005>

Received March 2023; Received in revised form October 2023; Accepted 20 November 2023

Available online 30 November 2023

0167-6687/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. Covariate partition for a Poisson-distributed simulation. Two covariates x_1, x_2 follow uniform distribution, i.e., $x_1, x_2 \sim U\{-3, -2, -1, 1, 2, 3\}$. The response variable, which is simulated for the points, has Poisson intensity equal to 1 (circles) and 7 (triangles). Each bar represents the average value (≈ 4) of Poisson intensity in that row/column of data.

techniques, random forests and boosting machines have been introduced in the context of insurance by adopting actuarial loss distributions in these models to capture the characteristics of insurance claims. We refer to Blier-Wong et al. (2020) for a recent literature review on this topic and Denuit and Trufin (2019); Wuthrich and Merz (2022); Wuthrich and Buser (2023) for more detailed discussion.

Insurance pricing models are heavily regulated and they must meet specific requirements before being deployed in practice, which posts some challenges for machine learning methods; see Henckaerts et al. (2021). Therein, it is stressed that pricing models must be transparent and easy to communicate to all the stakeholders and that the insurer has the social role of creating solidarity among the policy-holders so that the use of machine learning for pricing should in no way lead to an extreme penalization of risk or discrimination. The latter has also been noted recently in, e.g., Denuit et al. (2021); Wüthrich (2020) where it is claimed that prediction accuracy on an individual level should not be the ultimate goal in insurance pricing; one also needs to ensure the balance property. Bearing these points in mind, researchers have concluded that tree-based models are good candidates for insurance pricing (Henckaerts et al., 2021; Quan, 2019; Hu et al., 2022; Meng et al., 2022; Lindholm et al., 2022). More precisely, the use of CART, first introduced in Breiman et al. (1984), partitions a portfolio of policy-holders into smaller groups of homogeneous risk profiles based on some risk factors in which a constant prediction is used for each sub-group. This results in a highly transparent model and automatically induces solidarity among the policy-holders in a sub-group. Although a large number of scholars have carried out empirical and theoretical studies on the effectiveness of CART, limitations of the forward-search recursive partitioning method used in CART have been identified. In particular, the predictive performance tends to be low, and it is known to be unstable: small variations in the training set can result in greatly different trees and different predictions for the same test examples. Due to these limitations, more complex tree-based models that combine multiple trees in an ensemble have been popular in insurance prediction and pricing, but these ensemble techniques usually introduce additional difficulties in model transparency. In this paper and the sequel, we propose BCART models for insurance claims prediction. Instead of making an ensemble of trees, we look for one good tree, which can improve the prediction ability whilst ensuring model transparency, by adopting a Bayesian approach applied to CART.

BCART models were first introduced by Chipman et al. (1998) and Denison et al. (1998), independently. The method has two basic components, prior specification (for the tree and its terminal node parameters) and a stochastic search. The method is to obtain a posterior distribution given the prior, thus leading the stochastic search towards more promising tree models. Compared with the tree that CART generates by a greedy forward-search recursive partitioning method, the BCART model generates a much better tree by an effective Bayesian-motivated stochastic search algorithm. This has been justified by simulation examples (with Gaussian-distributed data) in the aforementioned papers. Here, we show another simulation example with Poisson-distributed data to illustrate the effectiveness of BCART. Specifically, we simulate 5,000 Poisson-distributed observations where the Poisson intensity depends on two explanatory variables (or covariates) x_1 and x_2 as illustrated in Fig. 1. (See also Subsection 4.2.1 for a slightly more general simulation example.) It is clear from the figure that the optimal partition of the covariate space consists of four regions where the data in each region should follow a homogeneous Poisson distribution. Note that the “standard” CART will not be able to find the correct partition of the data as the Poisson intensities are almost uniform for both marginal distributions (see Fig. 1) and no matter how the first split is chosen, it is difficult to distinguish different Poisson intensities on the resulting subsets. In contrast, the proposed Poisson BCART can retrieve the optimal tree structure since it has the ability to explore the tree space in a global way (for example, it can modify previously chosen splits).

Since BCART models and their ensemble version – the Bayesian Additive Regression Trees (BART) models – generally outperform other machine learning models, they have been extensively studied in the literature; see, e.g., Linero (2017); Chipman et al. (2010); Prado et al. (2021); Murray (2021); Hill et al. (2020) and references therein. In particular, their excellent empirical performance has also motivated works on their theoretical foundations; see Linero and Yang (2018); Rocková et al. (2020). However, in most of these studies, the focus has been on Gaussian-distributed data, with some exceptions such as Murray (2021); Linero et al. (2020). It turns out that a data augmentation approach is needed when dealing with general non-Gaussian data. The existing algorithms do not seem to be directly applicable to insurance data for prediction and pricing. To cover this gap, as a first step we propose BCART models for claims frequency taking account special features of insurance data such as the high number of zeros and involvement of exposures. We refer to Lee (2020, 2021) for a review of claims frequency modelling which also includes some nice analyses on exposures.

The main contributions of this paper are as follows:

- We give a general MCMC algorithm for the BCART models applied to any distributed data, where a data augmentation may be needed. In doing so, we follow some ideas in Meng and Van Dyk (1999); Van Dyk and Meng (2001).
- We introduce a novel model selection method for BCART models based on the deviance information criterion (DIC). Note that DIC was introduced in Spiegelhalter et al. (2002) which appeared a few years after the introduction of BCART (Chipman et al., 1998). The effectiveness of this approach is illustrated by several designed simulation examples and real insurance data.
- We implement the BCART for Poisson, NB and ZIP distributions which are not currently available in any existing R packages. In particular, we introduce two different ways of incorporating exposure in the NB and ZIP models, following the lines of study in Lee (2020, 2021). The simulation examples and real insurance data analysis show the applicability of these proposed BCART models.
- To date, Bayesian tree-based models have not attracted enough attention compared to other machine learning methods in the actuarial community. This first step of applying BCART for claims frequency modelling will open the door for more sophisticated tree-based models to meet the needs of the insurance industry.

Outline of the rest of the paper: In Section 2, we review the BCART framework which includes an extension with data augmentation and a model selection method using DIC. Section 3 introduces the notation for insurance claims frequency data and five BCART models including a Poisson model, two NB models and two ZIP models. In Section 4, we discuss the applicability of the proposed BCART models using three simulation examples and a real insurance claims dataset. Section 5 concludes the paper.

2. Bayesian CART

We shall briefly review the BCART framework of the seminal paper (Chipman et al., 1998). We begin with the general structure of a CART model. Consider a data set $(X, y) = ((x_1, y_1), (x_2, y_2), \dots, (x_n, y_n))^T$ with n observations. For the i -th observation, $x_i = (x_{i1}, x_{i2}, \dots, x_{ip})$ is a vector of p explanatory variables (or covariates) sampled from a space \mathcal{X} , while y_i is a response variable sampled from a space \mathcal{Y} . For our purpose of claims frequency modelling, \mathcal{Y} will be a set of non-negative integers.

A CART has two main components: a binary tree \mathcal{T} with b terminal nodes which induces a partition of the covariate space \mathcal{X} , denoted by $\{\mathcal{A}_1, \dots, \mathcal{A}_b\}$, and a parameter $\theta = (\theta_1, \theta_2, \dots, \theta_b)$ which associates the parameter value θ_t with the t -th terminal node. Note that here we do not specify the dimension and range of the parameter θ_t , which should be clear in the considered context below. If x_i is located in the t -th terminal node (i.e., $x_i \in \mathcal{A}_t$), then y_i has a distribution $f(y_i | \theta_t)$, where f represents a parametric family indexed by θ_t .

By associating observations with the b terminal nodes in the tree \mathcal{T} , we can represent the data set as

$$(X, y) = ((X_1, y_1), (X_2, y_2), \dots, (X_b, y_b))^T,$$

where $y_t = (y_{t1}, \dots, y_{tn_t})^T$ with n_t denoting the number of observations and y_{tj} denoting the j -th observation in the t -th terminal node, and X_t is an analogously defined $n_t \times p$ design matrix. We shall make the typical assumption that conditionally on (θ, \mathcal{T}) , response variables within a terminal node are independent and identically distributed (IID), and they are also independent across terminal nodes. The CART model likelihood in this case will take the form

$$p(y | X, \theta, \mathcal{T}) = \prod_{t=1}^b f(y_t | \theta_t) = \prod_{t=1}^b \prod_{i=1}^{n_t} f(y_{ti} | \theta_t). \tag{1}$$

It is worth noting that instead of the IID assumption within the terminal nodes more general models can be considered, see, e.g., Chipman et al. (2002, 2003) and the references therein.

Given that (θ, \mathcal{T}) determines a CART model, a Bayesian analysis of the problem is conducted by specifying a prior distribution $p(\theta, \mathcal{T})$, and inference about θ and \mathcal{T} will be based on the joint posterior $p(\theta, \mathcal{T} | y)$ using a suitable MCMC algorithm. Since θ indexes the parametric model whose dimension depends on the number of terminal nodes of the tree, it is usually convenient to apply the relationship

$$p(\theta, \mathcal{T}) = p(\theta | \mathcal{T})p(\mathcal{T}) \tag{2}$$

and specify the tree prior distribution $p(\mathcal{T})$ and the terminal node parameter prior distribution $p(\theta | \mathcal{T})$, respectively. This strategy, introduced by George (1998), offers several advantages for Bayesian model selection as outlined in Chipman et al. (1998).

2.1. Specification of tree prior $p(\mathcal{T})$

The prior for \mathcal{T} has two components: a tree topology and a decision rule for each of the internal/branch nodes. We shall adopt the branching process prior for the topology of \mathcal{T} proposed by Chipman et al. (1998). Due to its computational effectiveness using Metropolis-Hastings (MH) search algorithms, this prior specification has been the most popular in the literature. A draw from this prior is obtained by generating, for each node at depth d (with $d = 0$ for the root node), two child nodes with probability

$$p(d) = \gamma(1 + d)^{-\rho}, \tag{3}$$

where $\gamma > 0, \rho \geq 0$ are parameters controlling the structure and size of the tree. This process iterates for $d = 0, 1, \dots$, until we reach a depth at which all the nodes cease growing. Note that $p(d)$ is not a probability mass function, but instead is the probability of a given node at depth d being converted to a branch node. A sufficient condition for the termination of this branching process is that $\rho > 0$, and the case $\rho = 0$ corresponds to the Galton-Watson process, see, e.g., Athreya and Ney (2004). We refer to Linero (2018) for further theoretical discussion of this prior. Clearly, γ controls the overall rate of branching at a node, and the larger ρ becomes, the less likely that deeper nodes will branch, resulting in relatively smaller trees. In Chipman et al. (1998), some simulations about the number of terminal nodes associated with the values of the pair (γ, ρ) are carried out, which have been used as a guidance when choosing these parameters to generate trees with a certain number of terminal nodes.

After the tree topology is generated, each internal node is associated with a decision rule of the form $x_i < c_l$ or $x_i \in C_l$ according to whether x_i is a continuous or a categorical explanatory variable, where x_i is selected independently and uniformly among the available explanatory variables

for each internal node, and the split value c_l or split category subset C_l are selected uniformly among those available for the selected variable x_l . In practice, we only consider the overall set of possible split values to be finite; if the l -th variable is continuous, the grid for the variable is either uniformly spaced or given by a collection of observed quantiles of $\{x_{il}, i = 1, 2, \dots, n\}$. If the l -th variable is categorical, the split category subset C_l is usually selected uniformly among all possible subsets. However, this approach may not be efficient in the (Bayesian) tree search, particularly when the number of categorical levels of x_l is large. Instead, we shall adopt the same treatment of categorical variables as in the traditional CART greedy search algorithm. For example, in the Poisson case this is done as follows: calculate for each available categorical level, say k , of x_l in that node the empirical frequency $\bar{\lambda}_k(x_l)$ and use this empirical frequency $\bar{\lambda}_k(x_l)$ as a numerical replacement for the categorical level k of x_l . A subset C_l will be selected uniformly based on the ordered values $\bar{\lambda}_k(x_l)$.

Certainly, the design of tree prior can be more intricate than the one proposed in Chipman et al. (1998). There have been several alternatives discussed in the literature. In a recent contribution (Rocková et al., 2020), the convergence of the posterior distribution with a near-minimax concentration rate is studied, where it is shown that the original proposal given by (3) does not decay at a fast enough rate to guarantee the optimal rate of convergence. Instead, a sufficient condition for optimality is induced by the following probability

$$p(d) = \gamma^d, \quad \text{for some } 0 < \gamma < 1/2.$$

Most recently, it is noted in Saha (2023) that the original proposal (3) can still offer better empirical solutions. We believe further theoretical and empirical studies in this direction are still needed. An alternative to the branching process prior is to specify a prior directly on the number of leaves and a conditionally-uniform prior on the space of trees. In Denison et al. (1998), a Poisson-distributed prior is used for the number of leaves, and then a uniform prior over valid trees (i.e., trees with no empty bottom leaves) with that number of leaves is imposed. As noticed by Wu et al. (2007), the uniform prior over valid trees in Denison et al. (1998) tends to produce more unbalanced trees than balanced ones. Instead, they propose a pinball prior which can generate balanced or skewed trees by adjusting a hyper-parameter. Furthermore, instead of uniformly selecting the split value, a normal distribution is used for the split value in their simulation and real data analysis in Wu et al. (2007). Recently, some other tree priors have also been introduced for the purpose of variable selection (particularly when $p > n$), see, e.g., Bleich et al. (2014); Linero (2018); Rocková et al. (2020); Liu et al. (2021). In Linero (2018), the author proposes a sparsity-inducing Dirichlet prior for the splitting proportions of the explanatory variables, resulting in this prior allows the model to perform a fully Bayesian variable selection. Furthermore, in Rocková et al. (2020); Liu et al. (2021) a spike-and-tree variant is proposed by injecting one more layer on top of the prior used in Denison et al. (1998), that is, a prior over the active set of explanatory variables.

In our current implementation, we adopt the uniform specification for both variable and split value in each of the internal nodes, which is natural and simple. It is also noted in Chipman et al. (1998) that it would be beneficial to incorporate expert knowledge on the prior specification (i.e., using a non-uniform prior), however, our simulation studies in Section 4.2.1 show that using the uniform prior is able to identify the correct splitting rules even in the presence of noise variables. This seems to be a consequence of the Metropolis-Hastings random search steps, which tends to not accept noise splitting variables. We refer to Bleich et al. (2014) for some relevant discussions with the same conclusion.

2.2. Specification of the terminal node parameter prior $p(\theta | \mathcal{T})$

When choosing $p(\theta | \mathcal{T})$, it is vital to realize that employing priors that allow for analytical simplification can greatly reduce the computational burden of posterior calculation and exploration. This is especially true for the choice of the form $p(\theta | \mathcal{T})$ for which it is possible to analytically margin out θ to obtain the integrated likelihood

$$p(\mathbf{y} | \mathbf{X}, \mathcal{T}) = \int p(\mathbf{y} | \mathbf{X}, \theta, \mathcal{T}) p(\theta | \mathcal{T}) d\theta = \prod_{t=1}^b \int f(y_t | \theta_t) p(\theta_t) d\theta_t = \prod_{t=1}^b \int \prod_{i=1}^{n_t} f(y_{it} | \theta_t) p(\theta_t) d\theta_t, \tag{4}$$

where in the second equality we assume that conditional on the tree \mathcal{T} with b terminal nodes as above, the parameters $\theta_t, t = 1, 2, \dots, b$, have IID priors $p(\theta_t)$, which is a common assumption. Examples where this integration has a closed-form expression can be found in, e.g., Chipman et al. (1998); Linero (2017), particularly for Gaussian-distributed data \mathbf{y} . When no such priors can be found, we have to resort to the technique of data augmentation (see, e.g., Kindo et al. (2016); Linero et al. (2020); Murray (2021)) which will be discussed later. Combining the integrated likelihood $p(\mathbf{y} | \mathbf{X}, \mathcal{T})$ with tree prior $p(\mathcal{T})$, allows us to calculate the posterior of \mathcal{T}

$$p(\mathcal{T} | \mathbf{X}, \mathbf{y}) \propto p(\mathbf{y} | \mathbf{X}, \mathcal{T}) p(\mathcal{T}). \tag{5}$$

When using MCMC to conduct Bayesian inference, \mathcal{T} can be updated using an MH algorithm with the right-hand side of (5) used to compute the acceptance ratio. These MH simulations can be used to stochastically search the posterior space over trees to determine the high posterior probability trees from which we can choose a best one. The posterior sequence for θ is then obtained using an additional Gibbs sampler. It is worth noting that by integrating out θ in (4) we avoid the possible complexities associated with reversible jumps between continuous spaces of varying dimensions (Chipman et al., 2010; Green, 1995).

2.3. Stochastic search of posterior trees and parameters

Starting from the root node, the MCMC algorithm for simulating a Markov chain sequence of pairs $(\theta^{(1)}, \mathcal{T}^{(1)}), (\theta^{(2)}, \mathcal{T}^{(2)}), \dots$, using the posterior given in (5), is given in Algorithm 1.

In Algorithm 1, commonly used proposals (or transitions) for $q(\cdot, \cdot)$ include grow, prune, change and swap (see Chipman et al. (1998)), which are usually selected equal probability (i.e., 1/4 each). Other proposals have been suggested to improve the mixing of simulated trees, but these are often difficult to put into practice; see, e.g., Wu et al. (2007); Pratola (2016). One of the appealing features of these four proposals is that grow and prune steps are reversible counterparts of one another and both change and swap steps are independently reversible. As noticed in Chipman et al. (1998), this is very attractive for the calculation of $\alpha(\mathcal{T}^{(m)}, \mathcal{T}^*)$ in Algorithm 1, since there are substantial cancellations in the ratio (see also Kapelner and Bleich (2013) for detailed calculations). In our implementation, we consider these four proposals detailed as follows:

Algorithm 1 One step of the MCMC algorithm for updating the BCART parameterized by (θ, \mathcal{T}) .

Input: Data (X, y) and current values $(\theta^{(m)}, \mathcal{T}^{(m)})$
 1: Generate a candidate value \mathcal{T}^* with probability distribution $q(\mathcal{T}^{(m)}, \mathcal{T}^*)$
 2: Set the acceptance ratio $\alpha(\mathcal{T}^{(m)}, \mathcal{T}^*) = \min \left\{ \frac{q(\mathcal{T}^*, \mathcal{T}^{(m)}) p(y|X, \mathcal{T}^*)}{q(\mathcal{T}^{(m)}, \mathcal{T}^*) p(y|X, \mathcal{T}^{(m)})} \frac{p(\mathcal{T}^*)}{p(\mathcal{T}^{(m)})}, 1 \right\}$
 3: Update $\mathcal{T}^{(m+1)} = \mathcal{T}^*$ with probability $\alpha(\mathcal{T}^{(m)}, \mathcal{T}^*)$, otherwise, set $\mathcal{T}^{(m+1)} = \mathcal{T}^{(m)}$
 4: Sample $\theta^{(m+1)} \sim p(\theta | \mathcal{T}^{(m+1)}, X, y)$
Output: New values $(\theta^{(m+1)}, \mathcal{T}^{(m+1)})$

- **Grow:** Randomly select a terminal node. Split it into two new child nodes and randomly assign it a decision rule according to the prior specified in Section 2.1 until the resulting two child nodes satisfy a minimum observation requirement. If no such decision rule exists, draw a new terminal node (without replacement) and try again. If no such terminal node exists, stop grow.
 - **Prune:** A terminal node is randomly selected. The chosen node and its sibling node are pruned into the direct parent node which then becomes a new terminal node.
 - **Change:** apply one of the following two types of change to a selected internal node:
 - **Change1:** Reassign randomly only the split value/category subset according to the prior specified in Section 2.1.
 - **Change2:** Reassign randomly both the splitting variable and the corresponding split value/category subset according to the prior specified in Section 2.1.
- In each of the above changes, randomly select an internal node with the reassignment selected at random from a set (without replacement) until the updated nodes satisfy the minimum observation requirement. If no such reassignment exists, draw a new internal node (without replacement) and try again. If no such internal node exists, stop change.
- **Swap:** Randomly pick a parent-child pair which are both internal nodes and swap their decision rules until the updated nodes satisfy the minimum observation requirement. If no such parent-child pair exists, stop swap.

Remark 1. (a). Note that in step 4 of Algorithm 1, sampling of $\theta^{(m+1)}$ is needed only for those nodes that were involved in the proposed move from $\mathcal{T}^{(m)}$ to \mathcal{T}^* and only when this move was accepted.

(b). In comparison to Chipman et al. (1998), we apply two types of change moves as discussed in Denison et al. (1998). The introduction of these two types of change is helpful to improve the mixing of posterior trees, as demonstrated by our simulation study in Section 4.2.1. Moreover, it is noted that a swap between a parent-child pair with splits using the same variable is impossible. Considering this in our implementation improves the computational efficiency.

2.4. MCMC algorithm with data augmentation

In this section, we discuss the case where there is no obvious prior distribution $p(\theta_t)$ such that the integration in (4) is of closed-form, particularly, for non-Gaussian data y . In this case, we shall use a data augmentation method in implementing the MCMC algorithm. Some special cases have been discussed in Chipman et al. (2010); Kindo et al. (2016); Linero et al. (2020); Murray (2021).

The term data augmentation originated from Tanner and Wong’s data augmentation algorithm (Tanner and Wong, 1987). It is introduced purely for computational purposes and a latent variable is required so that the original distribution is the marginal distribution of the augmented one. We refer to Van Dyk and Meng (2001) for an overview of data augmentation and relevant theory. For our purpose, we augment the data y by introducing a latent variable $z = (z_1, z_2, \dots, z_n)$ so that the integration in (7) below is computable for augmented data (y, z) . To this end, we shall follow the idea of marginal augmentation introduced in Meng and Van Dyk (1999) (see also Van Dyk and Meng (2001)). In their framework, our parameter θ can be interpreted as a working parameter, and thus the integrated likelihood is given as

$$p(y | X, \mathcal{T}) = \int p(y, z | X, \mathcal{T}) dz, \tag{6}$$

where

$$p(y, z | X, \mathcal{T}) = \int p(y, z | X, \theta, \mathcal{T}) p(\theta | \mathcal{T}) d\theta = \prod_{t=1}^b \int f(y_t, z_t | \theta_t) p(\theta_t) d\theta_t = \prod_{t=1}^b \int \prod_{i=1}^{n_t} f(y_{ti}, z_{ti} | \theta_t) p(\theta_t) d\theta_t, \tag{7}$$

with $z_t = (z_{t1}, z_{t2}, \dots, z_{tn_t})$ defined according to the partition of \mathcal{X} and with obvious independence assumed. Following Scheme 3 of Meng and Van Dyk (1999) (see also Section 3 of Van Dyk and Meng (2001)), we propose the following Algorithm 2 to simulate a Markov chain sequence of pairs $(\theta^{(1)}, \mathcal{T}^{(1)}), (\theta^{(2)}, \mathcal{T}^{(2)}), \dots$, starting from the root node.

Algorithm 2 One step of the MCMC algorithm for updating the BCART parameterized by (θ, \mathcal{T}) using data augmentation.

Input: Data (X, y) and current values $(\theta^{(m)}, \mathcal{T}^{(m)}, z^{(m)})$
 1: Generate a candidate value \mathcal{T}^* with probability distribution $q(\mathcal{T}^{(m)}, \mathcal{T}^*)$
 2: Sample $z^{(m+1)} \sim p(z | X, y, \theta^{(m)}, \mathcal{T}^{(m)})$
 3: Set the acceptance ratio $\alpha(\mathcal{T}^{(m)}, \mathcal{T}^*) = \min \left\{ \frac{q(\mathcal{T}^*, \mathcal{T}^{(m)}) p(y, z^{(m+1)} | X, \mathcal{T}^*)}{q(\mathcal{T}^{(m)}, \mathcal{T}^*) p(y, z^{(m+1)} | X, \mathcal{T}^{(m)})} \frac{p(\mathcal{T}^*)}{p(\mathcal{T}^{(m)})}, 1 \right\}$
 4: Update $\mathcal{T}^{(m+1)} = \mathcal{T}^*$ with probability $\alpha(\mathcal{T}^{(m)}, \mathcal{T}^*)$, otherwise, set $\mathcal{T}^{(m+1)} = \mathcal{T}^{(m)}$
 5: Sample $\theta^{(m+1)} \sim p(\theta | \mathcal{T}^{(m+1)}, X, y, z^{(m+1)})$
Output: New values $(\theta^{(m+1)}, \mathcal{T}^{(m+1)}, z^{(m+1)})$

Note that in some cases introducing one latent variable z is insufficient to obtain a closed-form for the integration in (7); more latent variables may be required. In that case, we can easily extend Algorithm 2 to include multivariate latent variables and use the Gibbs sampler in step 2. Clearly, the more latent variables used, the slower the convergence of the Markov chain sequence. As discussed in Van Dyk and Meng (2001), it is an “art” to search for efficient data augmentation schemes. We discuss this point later for the claims frequency models.

Remark 2. Similar to Algorithm 1, in step 2 and step 5 of Algorithm 2 the sampling is needed only for those nodes that were involved in the proposed move from $\mathcal{T}^{(m)}$ to \mathcal{T}^* , and step 5 is needed only when this move was accepted.

2.5. Posterior tree selection and prediction

The MCMC algorithms described in the previous section can be used to search for desirable trees. However, as discussed in Chipman et al. (1998) and illustrated below in our analysis, the algorithms quickly converge and then move locally in that region for a long time, which occurs because proposals make local moves over a sharply peaked multimodal posterior. Instead of making long runs of search to move from one mode to another better one, we follow the idea of Chipman et al. (1998) to repeatedly restart the algorithm. As many trees are visited by each run of the algorithm, we need a method to identify those trees which are of most interest. Moreover, the structure of trees in the convergence regions is mostly determined by the hyper-parameters γ, ρ which also need to be chosen appropriately. In Chipman et al. (1998), the integrated likelihood $p(\mathbf{y} | \mathbf{X}, T)$ is used as a measure to choose good trees from one run of the algorithm, though other measures, like residual sum of squares, could also be introduced. However, there is no discussion on how the tree prior hyper-parameters γ, ρ should be determined optimally. A natural way to deal with this is to use cross-validation which, however, requires repeated model fits and is very computationally expensive. In this paper, we propose to use DIC for choosing appropriate γ, ρ , and thus introduce a three-step approach for selecting an “optimal” tree among those visited. To this end, we first give a definition of DIC for a Bayesian CART. We refer to Spiegelhalter et al. (2002); Celeux et al. (2006); Gelman et al. (2014); Spiegelhalter et al. (2014) for more detailed discussion of DIC and its extensions.

Consider the tree \mathcal{T} with b terminal nodes and parameters $\theta_t, t = 1, 2, \dots, b$, previously defined. We first introduce DIC for each node using the standard definition, the DIC for the tree is then defined as the sum of the DIC of all terminal nodes in the tree due to the independence assumption. For node t , we call

$$D(\theta_t) = -2 \log(f(\mathbf{y}_t | \theta_t)) = -2 \sum_{i=1}^{n_t} \log(f(y_{ti} | \theta_t)) \tag{8}$$

the *deviance*.

Analogously to Akaike’s information criterion (AIC), Spiegelhalter et al. (2002) proposed the DIC based on the principle DIC=“goodness of fit”+“complexity”, which is defined as

$$DIC_t = D(\bar{\theta}_t) + 2p_{D_t},$$

where $\bar{\theta}_t = E_{\text{post}}(\theta_t)$ is the posterior mean (with E_{post} denoting expectation over the posterior distribution of θ given data \mathbf{y}), and p_{D_t} is the *effective number of parameters* given by

$$p_{D_t} = \overline{D(\theta_t)} - D(\bar{\theta}_t) = -2E_{\text{post}}(\log(f(\mathbf{y}_t | \theta_t))) + 2 \log(f(\mathbf{y}_t | \bar{\theta}_t)) = 2 \sum_{i=1}^{n_t} \left(\log(f(y_{ti} | \bar{\theta}_t)) - E_{\text{post}}(\log(f(y_{ti} | \theta_t))) \right). \tag{9}$$

The DIC of the tree \mathcal{T} with b terminal nodes is then defined as

$$DIC := \sum_{t=1}^b DIC_t = D(\bar{\theta}) + 2p_D, \tag{10}$$

where $D(\bar{\theta}) = \sum_{t=1}^b D(\bar{\theta}_t)$ and $p_D = \sum_{t=1}^b p_{D_t}$ are the deviance and effective number of parameters of the tree.

Next, we introduce DIC for tree models with data augmentation. Depending on whether the latent variable \mathbf{z} is treated as a parameter or not, there are three types of likelihoods leading to eight versions of DIC as discussed in Celeux et al. (2006). Due to the complexity in implementing any of those eight and motivated by the idea that DIC=“goodness of fit”+“complexity”, we introduce a new DIC for node t in the tree as follows

$$DIC_t = D(\bar{\theta}_t) + 2q_{D_t}, \tag{11}$$

where $D(\bar{\theta}_t)$ is the deviance defined through the data \mathbf{y}_t (as in (8)) which represents the goodness of fit, and q_{D_t} is the *effective number of parameters* defined through the augmented data $(\mathbf{y}_t, \mathbf{z}_t)$ as follows

$$q_{D_t} = -2E_{\text{post}}(\log(f(\mathbf{y}_t, \mathbf{z}_t | \theta_t))) + 2 \log(f(\mathbf{y}_t, \mathbf{z}_t | \bar{\theta}_t)) = 2 \sum_{i=1}^{n_t} \left(\log(f(y_{ti}, z_{ti} | \bar{\theta}_t)) - E_{\text{post}}(\log(f(y_{ti}, z_{ti} | \theta_t))) \right), \tag{12}$$

where $\bar{\theta}_t = E_{\text{post}}(\theta_t)$, and in this case E_{post} denotes expectation over the posterior distribution of θ given augmented data (\mathbf{y}, \mathbf{z}) . As we will see below, for the frequency models, q_{D_t} is approximately the dimension of θ_t as the sample size n_t in node t tends to infinity. Similarly, the DIC of tree \mathcal{T} with b terminal nodes is thus defined as

$$DIC = D(\bar{\theta}) + 2q_D, \tag{13}$$

where $q_D = \sum_{t=1}^b q_{D_t}$.

Remark 3. (a). Note that DIC is defined using plug-in prediction densities $f(y_{ii} | \bar{\theta}_t)$ in (9) (similarly $f(y_{ii}, z_{ii} | \bar{\theta}_t)$ in (12)). More recently, a new criterion called WAIC was introduced by Watanabe and Opper (2010) (see also Gelman et al. (2014); Spiegelhalter et al. (2014)), where in its definition the plug-in prediction density is replaced by the full prediction density $E_{\text{post}}(f(y_{ii} | \theta_t))$. When the explicit expression is not available, this posterior expectation is usually computed by a Monte Carlo algorithm as $S^{-1} \sum_{k=1}^S f(y_{ii} | \theta^k)$, where θ^k is simulated from the posterior distribution of θ_t . In the following section, we will see that this posterior expectation can be obtained explicitly for the Poisson model, but not for other models. It turns out that using WAIC gives the same selected model as DIC in our simulation examples. Additionally, since it involves Monte Carlo algorithm and as such could be considerably more computationally expensive, we suggest using DIC.

Table 1
Three-step approach for “optimal” tree selection.

Step 1:	Set a sequence of hyper-parameters $(\gamma_j, \rho_j), j = m_s, \dots, m_e$, such that for (γ_j, ρ_j) , the MCMC algorithm converges to a region of trees which have j terminal nodes.
Step 2:	For each j in Step 1, select the tree with maximum likelihood $p(\mathbf{y} \mathbf{X}, \bar{\theta}, \mathcal{T})$ from the convergence region.
Step 3:	From the trees obtained in Step 2, select the optimal one using DIC.

(b). It is worth noting that if the independence assumption within the terminal nodes is violated (e.g., Chipman et al. (2002, 2003)), the DIC may also be used as a tool for model selection but the formulation would not be of the simple summation form as in (8). We refer to Spiegelhalter et al. (2002) for examples and relevant discussions.

Now, we are ready to introduce the three-step approach for selecting an “optimal” tree from the MCMC algorithms. Let $m_s < m_e$ be two user input integers which represent the belief that the optimal number of terminal nodes lies in $[m_s, m_e]$. In practice, these can be estimated first by using some other methods, e.g., a standard CART model. The three-step approach is described in Table 1. In what follows, the tree selected by using the three-step approach will be called an “optimal” tree.

Remark 4. (a). The relation between hyper-parameters (γ_j, ρ_j) and the distribution of the number of terminal nodes of tree has been illustrated in Chipman et al. (1998). It does not seem hard to set values for (γ_j, ρ_j) so that the MCMC algorithms will converge to a region of trees with required j terminal nodes. It is also worth noting that the distribution of the number of terminal nodes is also affected by the data in hand, which can be seen from the calculation of the acceptance ratio in the MCMC algorithms. In our simulations and real data analysis below, we have to select a relatively larger ρ in order to achieve our goals.

(b). In Step 2, the so-called data likelihood $p(\mathbf{y} | \mathbf{X}, \bar{\theta}, \mathcal{T})$, rather than the integrated likelihood $p(\mathbf{y} | \mathbf{X}, \mathcal{T})$, is used, which is due to our interest in the fit of the parametric model to data. The simulations and real data in Section 4 indicate that these two types of likelihood show a consistency in the ordering of their values, and thus we suspect there is no big difference using either of them.

Suppose \mathcal{T} with b terminal nodes and parameter $\bar{\theta}$ is the optimal tree obtained from the above three-step approach. For a given new \mathbf{x} the predicted \hat{y} using this tree model is defined as

$$\hat{y} | \mathbf{x} = \sum_{t=1}^b E(y | \bar{\theta}_t) I_{(\mathbf{x} \in \mathcal{A}_t)}, \tag{14}$$

where $I_{(\cdot)}$ denotes the indicator function and $\{\mathcal{A}_t\}_{t=1}^b$ is the partition of \mathcal{X} by \mathcal{T} .

Remark 5. An alternative prediction given \mathbf{x} can be defined using the full predictive density as

$$\hat{y} | \mathbf{x} = \sum_{t=1}^b E_{\text{post}}(E(y | \theta_t)) I_{(\mathbf{x} \in \mathcal{A}_t)}. \tag{15}$$

However, for the frequency models the explicit expression can be found only for the Poisson case, and for other models the Monte Carlo method is needed to estimate the posterior expectation. Thus, we shall use (14) for simplicity.

3. Bayesian CART claims frequency models

In this section, we introduce the BCART for insurance claims frequency by specifying the response distribution in the general framework introduced in Section 2. We shall discuss three commonly used distributions in the literature to model the claim numbers, namely, Poisson, NB and ZIP distributions; see, e.g., Wuthrich and Merz (2022); Lee (2020, 2021). To this end, we first introduce the claims data. A claims data set with n policy-holders can be described by $(\mathbf{X}, \mathbf{v}, \mathbf{N}) = ((\mathbf{x}_1, v_1, N_1), \dots, (\mathbf{x}_n, v_n, N_n))^T$, where $\mathbf{x}_i = (x_{i1}, \dots, x_{ip}) \in \mathcal{X}$ consists of rating variables (e.g., area, driver age, car brand in car insurance); N_i is the number of claims reported, and $v_i \in (0, 1]$ is the exposure in yearly units which is used to quantify how long the policy-holder is exposed to risk. The goal is to explain and predict the claims information N_i based on the rating variables \mathbf{x}_i and the exposure v_i for each individual policy i , which leads to the *claims frequency*, i.e., the number of claims filed per unit year of exposure to risk. We will discuss below how this can be done with BCART models.

3.1. Poisson model

Consider a tree \mathcal{T} with b terminal nodes as discussed in Section 2. In a Poisson model, we assume

$$N_i | \mathbf{x}_i, v_i \sim \text{Poi} \left(v_i \sum_{t=1}^b \lambda_t I_{(\mathbf{x}_i \in \mathcal{A}_t)} \right)$$

for the i -th observation where \mathcal{A}_t is a partition of \mathcal{X} . Here we use the standard notation λ_t for claims frequency rather than the generic notation θ_t for the parameter in terminal node t . Essentially, we have specified the distribution $f(y_i | \theta_t)$ for terminal node t (see Section 2) as

$$f_P(m | \lambda_t) = P(N_i = m | \lambda_t) = \frac{e^{-\lambda_t v_i} (\lambda_t v_i)^m}{m!}, \quad m = 0, 1, 2, \dots, \tag{16}$$

for the i -th observation such that $\mathbf{x}_i \in \mathcal{A}_t$. Note that, for simplicity, here and hereafter, the exposure v_i and \mathbf{x}_i will be compressed in some notation. Based on the discussions in Section 2.2, we choose the gamma prior for λ_t with hyper-parameters $\alpha, \beta > 0$, that is,

$$p(\lambda_t) = \frac{\beta^\alpha \lambda_t^{\alpha-1} e^{-\beta \lambda_t}}{\Gamma(\alpha)}, \tag{17}$$

with $\Gamma(\cdot)$ denoting the gamma function. As in Section 2, for terminal node t we define the associated data as $(\mathbf{X}_t, \mathbf{v}_t, \mathbf{N}_t) = ((X_{t1}, v_{t1}, N_{t1}), \dots, (X_{tn_t}, v_{tn_t}, N_{tn_t}))^\top$. With the above gamma prior, the integrated likelihood for terminal node t can be obtained as

$$\begin{aligned} p_P(\mathbf{N}_t | \mathbf{X}_t, \mathbf{v}_t) &= \int_0^\infty f_P(\mathbf{N}_t | \lambda_t) p(\lambda_t) d\lambda_t \\ &= \int_0^\infty \prod_{i=1}^{n_t} \frac{e^{-\lambda_t v_{ti}} (\lambda_t v_{ti})^{N_{ti}}}{N_{ti}!} \frac{\beta^\alpha \lambda_t^{\alpha-1} e^{-\beta \lambda_t}}{\Gamma(\alpha)} d\lambda_t \\ &= \frac{\beta^\alpha \prod_{i=1}^{n_t} v_{ti}^{N_{ti}}}{\Gamma(\alpha) \prod_{i=1}^{n_t} N_{ti}!} \int_0^\infty \lambda_t^{\sum_{i=1}^{n_t} N_{ti} + \alpha - 1} e^{-(\sum_{i=1}^{n_t} v_{ti} + \beta)\lambda_t} d\lambda_t \\ &= \frac{\beta^\alpha \prod_{i=1}^{n_t} v_{ti}^{N_{ti}}}{\Gamma(\alpha) \prod_{i=1}^{n_t} N_{ti}!} \frac{\Gamma(\sum_{i=1}^{n_t} N_{ti} + \alpha)}{(\sum_{i=1}^{n_t} v_{ti} + \beta)^{\sum_{i=1}^{n_t} N_{ti} + \alpha}}. \end{aligned} \tag{18}$$

Clearly, from (18), we see that the posterior distribution of λ_t , conditional on \mathbf{N}_t , is given by

$$\lambda_t | \mathbf{N}_t \sim \text{Gamma}\left(\sum_{i=1}^{n_t} N_{ti} + \alpha, \sum_{i=1}^{n_t} v_{ti} + \beta\right). \tag{19}$$

The integrated likelihood for the tree \mathcal{T} is thus given by

$$p_P(\mathbf{N} | \mathbf{X}, \mathbf{v}, \mathcal{T}) = \prod_{t=1}^b p_P(\mathbf{N}_t | \mathbf{X}_t, \mathbf{v}_t). \tag{20}$$

Next, we discuss the DIC for this tree, focusing on DIC_t for terminal node t . First, we have

$$D(\lambda_t) = -2 \sum_{i=1}^{n_t} \log f_P(N_{ti} | \lambda_t) = -2 \sum_{i=1}^{n_t} (-\lambda_t v_{ti} + N_{ti} \log(\lambda_t v_{ti}) - \log(N_{ti}!)), \tag{21}$$

and by (19) we get the posterior mean for λ_t as

$$\bar{\lambda}_t = E_{\text{post}}(\lambda_t) = \frac{\sum_{i=1}^{n_t} N_{ti} + \alpha}{\sum_{i=1}^{n_t} v_{ti} + \beta}. \tag{22}$$

Furthermore, we derive that

$$\begin{aligned} \overline{D(\lambda_t)} &= E_{\text{post}}(D(\lambda_t)) \\ &= 2 \sum_{i=1}^{n_t} v_{ti} E_{\text{post}}(\lambda_t) - 2 \sum_{i=1}^{n_t} N_{ti} E_{\text{post}}(\log(\lambda_t) + \log(v_{ti})) + 2 \sum_{i=1}^{n_t} \log(N_{ti}!) \\ &= 2 \left(\frac{\sum_{i=1}^{n_t} N_{ti} + \alpha}{\sum_{i=1}^{n_t} v_{ti} + \beta} \right) \sum_{i=1}^{n_t} v_{ti} - 2 \left(\psi \left(\sum_{i=1}^{n_t} N_{ti} + \alpha \right) - \log \left(\sum_{i=1}^{n_t} v_{ti} + \beta \right) \right) \sum_{i=1}^{n_t} N_{ti} - 2 \sum_{i=1}^{n_t} N_{ti} \log(v_{ti}) + 2 \sum_{i=1}^{n_t} \log(N_{ti}!), \end{aligned} \tag{23}$$

where we have used the fact that

$$E_{\text{post}}(\log(\lambda_t)) = \psi \left(\sum_{i=1}^{n_t} N_{ti} + \alpha \right) - \log \left(\sum_{i=1}^{n_t} v_{ti} + \beta \right),$$

with $\psi(x) = \Gamma'(x)/\Gamma(x)$ being the digamma function. Using (21)–(23), we obtain the effective number of parameters for terminal node t as

$$p_{D_t} = \overline{D(\lambda_t)} - D(\bar{\lambda}_t) = 2 \left(\log \left(\sum_{i=1}^{n_t} N_{ti} + \alpha \right) - \psi \left(\sum_{i=1}^{n_t} N_{ti} + \alpha \right) \right) \sum_{i=1}^{n_t} N_{ti},$$

and

$$\begin{aligned} \text{DIC}_t &= D(\bar{\lambda}_t) + 2p_{D_t} \\ &= 2 \left(\frac{\sum_{i=1}^{n_t} N_{ti} + \alpha}{\sum_{i=1}^{n_t} v_{ti} + \beta} \right) \sum_{i=1}^{n_t} v_{ti} - 2 \sum_{i=1}^{n_t} N_{ti} \left(\log \left(\frac{\sum_{i=1}^{n_t} N_{ti} + \alpha}{\sum_{i=1}^{n_t} v_{ti} + \beta} \right) + \log(v_{ti}) \right) + 2 \sum_{i=1}^{n_t} \log(N_{ti}!) \\ &\quad + 4 \left(\log \left(\sum_{i=1}^{n_t} N_{ti} + \alpha \right) - \psi \left(\sum_{i=1}^{n_t} N_{ti} + \alpha \right) \right) \sum_{i=1}^{n_t} N_{ti}. \end{aligned}$$

Then the DIC of tree \mathcal{T} is obtained using (10).

Remark 6. Since $\psi(x) = \log(x) - \frac{1}{2x}(1 + o(x))$, as $x \rightarrow \infty$, we immediately see that $p_{D_t} \rightarrow 1$ as $n_t \rightarrow \infty$. This explains the name of effective number of parameters in the Bayesian framework, as 1 is the number of parameters in the terminal node t for Poisson model if a flat prior is assumed for λ_t .

With the above (19)–(20) and DIC obtained, we can use the three-step approach proposed in Section 2.5 to search for an optimal tree, where (19) and (20) should be used in step 4 and step 2, respectively, in Algorithm 1. Given an optimal tree, the estimated claims frequency $\bar{\lambda}_t$ in terminal node t can be given by the posterior mean in (22), using (14). It is worth noting that we can obtain the same estimate by using (15) instead.

3.2. Negative binomial models

The NB distribution, a member of mixed Poisson family, offers an effective way to handle over-dispersed insurance claims frequency data where excessive zeros are common.

Consider the tree \mathcal{T} with b terminal nodes as before. In the NB model, we assume that $N_{ii} | X_{ii}, v_{ii}$ follows a NB distribution for all terminal nodes, $t = 1, \dots, b$. There are different ways to parameterize the NB distribution, particularly with the exposure, see, e.g., Lee (2020); Wuthrich and Merz (2022). We shall discuss two models below.

3.2.1. Negative binomial model 1 (NB1)

We first adopt the most common parameterization of the NB distribution, see, e.g., Murray (2021). That is, for terminal node t ,

$$f_{\text{NB1}}(m | \kappa_t, \lambda_t) = P(N_{ii} = m | \kappa_t, \lambda_t) = \frac{\Gamma(m + \kappa_t)}{\Gamma(\kappa_t)m!} \left(\frac{\kappa_t}{\kappa_t + \lambda_t v_{ii}}\right)^{\kappa_t} \left(\frac{\lambda_t v_{ii}}{\kappa_t + \lambda_t v_{ii}}\right)^m, \quad m = 0, 1, \dots, \tag{24}$$

where $\kappa_t, \lambda_t > 0$. It is easy to show that the mean and variance of N_{ii} are given by

$$E(N_{ii} | \kappa_t, \lambda_t) = \lambda_t v_{ii}, \quad \text{Var}(N_{ii} | \kappa_t, \lambda_t) = \lambda_t v_{ii} \left(1 + \frac{\lambda_t v_{ii}}{\kappa_t}\right). \tag{25}$$

The degree of over-dispersion in relation to the Poisson is controlled by the additional parameter κ_t in the NB model, which converges to the Poisson model as $\kappa_t \rightarrow \infty$.

In NB regression, the lack of simple and efficient algorithms for posterior computation has seriously limited routine applications of Bayesian approaches. Recent studies make Bayesian approaches appealing by introducing data augmentation techniques; see, e.g., Zhou et al. (2012); Murray (2021). In order to save on total computational time of the algorithm and avoid the difficulty of finding an appropriate prior for κ_t with corresponding data augmentation, we shall treat the parameter κ_t as known in the Bayesian framework which can be estimated upfront by using, e.g., the moment matching method. However, in line with the Poisson model, we shall treat λ_t as uncertain and use a gamma prior with corresponding data augmentation. Based on the formulas given in (25), we can estimate the parameter κ_t , using the moment matching method, see, e.g., Chapter 2 of Wuthrich (2022) as follows

$$\hat{\kappa}_t = \frac{\hat{\lambda}_t^2}{\hat{v}_t^2 - \hat{\lambda}_t} \frac{1}{n_t - 1} \left(\sum_{i=1}^{n_t} v_{ii} - \frac{\sum_{i=1}^{n_t} v_{ii}^2}{\sum_{i=1}^{n_t} v_{ii}} \right), \tag{26}$$

where

$$\hat{v}_t^2 = \frac{1}{n_t - 1} \sum_{i=1}^{n_t} v_{ii} \left(\frac{N_{ii}}{v_{ii}} - \hat{\lambda}_t \right)^2, \quad \hat{\lambda}_t = \frac{\sum_{i=1}^{n_t} N_{ii}}{\sum_{i=1}^{n_t} v_{ii}}. \tag{27}$$

Next, introducing a latent variable $\xi_t = (\xi_{t1}, \xi_{t2}, \dots, \xi_{tn_t}) \in (0, \infty)^{n_t}$, we can define a data augmented likelihood for the i -th data instance in terminal node t as

$$f_{\text{NB1}}(N_{ii}, \xi_{ii} | \hat{\kappa}_t, \lambda_t) = (\lambda_t v_{ii})^{N_{ii}} e^{-\xi_{ii} \lambda_t v_{ii}} \frac{\hat{\kappa}_t^{\hat{\kappa}_t} \xi_{ii}^{\hat{\kappa}_t + N_{ii} - 1} e^{-\xi_{ii} \hat{\kappa}_t}}{\Gamma(\hat{\kappa}_t) N_{ii}!}. \tag{28}$$

It is easily checked that integrating over $\xi_{ii} \in (0, \infty)$ in (28) yields the marginal distribution (24). Further, we see that ξ_{ii} , given data N_{ii} and parameters, is gamma distributed, i.e.,

$$\xi_{ii} | N_{ii}, \hat{\kappa}_t, \lambda_t \sim \text{Gamma}(\hat{\kappa}_t + N_{ii}, \hat{\kappa}_t + \lambda_t v_{ii}). \tag{29}$$

Given the data augmented likelihood in (28), the estimated parameter $\hat{\kappa}_t$ using (26), and a conjugate gamma prior for λ_t with hyper-parameters $\alpha, \beta > 0$ (cf. (17)), we can derive the integrated augmented likelihood for the terminal node t as follows

$$\begin{aligned} p_{\text{NB1}}(N_t, \xi_t | X_t, v_t, \hat{\kappa}_t) &= \int_0^\infty f_{\text{NB1}}(N_t, \xi_t | \hat{\kappa}_t, \lambda_t) p(\lambda_t) d\lambda_t \\ &= \int_0^\infty \prod_{i=1}^{n_t} \left[(\lambda_t v_{ii})^{N_{ii}} e^{-\xi_{ii} \lambda_t v_{ii}} \frac{\hat{\kappa}_t^{\hat{\kappa}_t} \xi_{ii}^{\hat{\kappa}_t + N_{ii} - 1} e^{-\xi_{ii} \hat{\kappa}_t}}{\Gamma(\hat{\kappa}_t) N_{ii}!} \right] \frac{\beta^\alpha \lambda_t^{\alpha-1} e^{-\beta \lambda_t}}{\Gamma(\alpha)} d\lambda_t \\ &= \frac{\hat{\kappa}_t^{\hat{\kappa}_t} \beta^\alpha}{\Gamma(\hat{\kappa}_t) \Gamma(\alpha)} \prod_{i=1}^{n_t} \left[\frac{v_{ii}^{N_{ii}}}{N_{ii}!} \xi_{ii}^{\hat{\kappa}_t + N_{ii} - 1} e^{-\xi_{ii} \hat{\kappa}_t} \right] \frac{\Gamma(\sum_{i=1}^{n_t} N_{ii} + \alpha)}{(\sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta)^{\sum_{i=1}^{n_t} N_{ii} + \alpha}}. \end{aligned} \tag{30}$$

Moreover, from the above we see that the posterior distribution of λ_t , given the augmented data (N_t, ξ_t) , is given by

$$\lambda_t \mid N_t, \xi_t \sim \text{Gamma} \left(\sum_{i=1}^{n_t} N_{ii} + \alpha, \sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta \right).$$

The integrated augmented likelihood for the tree \mathcal{T} is thus given by

$$p_{\text{NB1}}(N, \xi \mid X, v, \hat{\kappa}, \mathcal{T}) = \prod_{t=1}^b p_{\text{NB1}}(N_t, \xi_t \mid X_t, v_t, \hat{\kappa}_t). \tag{31}$$

Now, we discuss the DIC for this tree. Since we only consider uncertainty for λ but not for κ , the DIC defined in (13) cannot be adopted directly. Thus, using the idea that DIC=“goodness of fit”+“complexity”, we can introduce a new DIC_t for terminal node t as follows

$$\text{DIC}_t = D(\bar{\lambda}_t) + 2r_{D_t}.$$

Here, the goodness of fit is given by

$$D(\bar{\lambda}_t) = -2 \sum_{i=1}^{n_t} \log f_{\text{NB1}}(N_{ii} \mid \hat{\kappa}_t, \bar{\lambda}_t),$$

and the effective number of parameters r_{D_t} is given by

$$r_{D_t} = 1 + 2 \sum_{i=1}^{n_t} \left(\log(f_{\text{NB1}}(N_{ii}, \xi_{ii} \mid \hat{\kappa}_t, \bar{\lambda}_t)) - E_{\text{post}}(\log(f_{\text{NB1}}(N_{ii}, \xi_{ii} \mid \hat{\kappa}_t, \lambda_t))) \right), \tag{32}$$

where 1 represents the number for κ_t and the second part is for λ_t ,

$$\bar{\lambda}_t = \frac{\sum_{i=1}^{n_t} N_{ii} + \alpha}{\sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta},$$

and

$$\begin{aligned} & E_{\text{post}}(\log(f_{\text{NB1}}(N_{ii}, \xi_{ii} \mid \hat{\kappa}_t, \lambda_t))) \\ &= -2 \sum_{i=1}^{n_t} N_{ii} \left(\log(v_{ii}) + \psi \left(\sum_{i=1}^{n_t} N_{ii} + \alpha \right) - \log \left(\sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta \right) \right) + 2 \left(\frac{\sum_{i=1}^{n_t} N_{ii} + \alpha}{\sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta} \right) \sum_{i=1}^{n_t} \xi_{ii} v_{ii} \\ &+ 2 \sum_{i=1}^{n_t} (\log(N_{ii}!) - (\hat{\kappa}_t + N_{ii} - 1) \log(\xi_{ii} + \xi_{ii} \hat{\kappa}_t) + 2(\log(\Gamma(\hat{\kappa}_t)) - \hat{\kappa}_t \log(\hat{\kappa}_t))). \end{aligned} \tag{33}$$

Therefore, a direct calculation shows that the effective number of parameters for terminal node t is given by

$$r_{D_t} = 1 + 2 \left(\log \left(\sum_{i=1}^{n_t} N_{ii} + \alpha \right) - \psi \left(\sum_{i=1}^{n_t} N_{ii} + \alpha \right) \right) \sum_{i=1}^{n_t} N_{ii},$$

and thus

$$\begin{aligned} \text{DIC}_t &= D(\bar{\lambda}_t) + 2r_{D_t} \\ &= -2 \sum_{i=1}^{n_t} N_{ii} \left(\log(v_{ii}) + \log \left(\sum_{i=1}^{n_t} N_{ii} + \alpha \right) - \log \left(\sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta \right) \right) + 2 \left(\frac{\sum_{i=1}^{n_t} N_{ii} + \alpha}{\sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta} \right) \sum_{i=1}^{n_t} \xi_{ii} v_{ii} \\ &+ 2 \sum_{i=1}^{n_t} (\log(N_{ii}!) - (\hat{\kappa}_t + N_{ii} - 1) \log(\xi_{ii} + \xi_{ii} \hat{\kappa}_t) + 2(\log(\Gamma(\hat{\kappa}_t)) - \hat{\kappa}_t \log(\hat{\kappa}_t))) \\ &+ 2 + 4 \left(\log \left(\sum_{i=1}^{n_t} N_{ii} + \alpha \right) - \psi \left(\sum_{i=1}^{n_t} N_{ii} + \alpha \right) \right) \sum_{i=1}^{n_t} N_{ii}. \end{aligned}$$

3.2.2. Negative binomial model 2 (NB2)

We now consider another parameterization of the NB distribution, see, e.g., Lee (2020); Wuthrich and Merz (2022). Now, for terminal node t ,

$$f_{\text{NB2}}(m \mid \kappa_t, \lambda_t) = P(N_{ii} = m \mid \kappa_t, \lambda_t) = \frac{\Gamma(m + \kappa_t v_{ii})}{\Gamma(\kappa_t v_{ii}) m!} \left(\frac{\kappa_t}{\kappa_t + \lambda_t} \right)^{\kappa_t v_{ii}} \left(\frac{\lambda_t}{\kappa_t + \lambda_t} \right)^m, \quad m = 0, 1, \dots \tag{34}$$

It is easy to show that the mean of N_{ii} is the same as in (25), but the variance becomes

$$\text{Var}(N_{ii} \mid \kappa_t, \lambda_t) = \lambda_t v_{ii} \left(1 + \frac{\lambda_t}{\kappa_t} \right). \tag{35}$$

This formulation yields a fixed over-dispersion of size λ_t/κ_t which does not depend on the exposure v_{ii} , and thus it is sometimes preferred (see Wuthrich and Merz (2022)) and has been judged as more effective for real insurance data analysis (see Lee (2020)).

We use the same way to deal with κ_t and λ_t as in the previous subsection. Using the same approach as Chapter 2 of Wuthrich (2022), we can estimate the parameter κ_t as follows

$$\hat{\kappa}_t = \frac{\hat{\lambda}_t^2}{\hat{v}_t^2 - \hat{\lambda}_t}, \tag{36}$$

where \hat{v}_t^2 and $\hat{\lambda}_t$ are given in (27). Note that this parameterization offers a simpler estimation for $\hat{\kappa}_t$, and that $\hat{\lambda}_t$ is a minimal variance estimator; see Wuthrich (2022).

Similarly as before, we can define a data augmented likelihood for the i -th data instance in terminal node t as

$$f_{\text{NB2}}(N_{ii}, \xi_{ii} | \hat{\kappa}_t, \lambda_t) = (\lambda_t v_{ii})^{N_{ii}} e^{-\xi_{ii} \lambda_t v_{ii}} \frac{(\hat{\kappa}_t v_{ii})^{\hat{\kappa}_t v_{ii}} \xi_{ii}^{\hat{\kappa}_t v_{ii} + N_{ii} - 1} e^{-\xi_{ii} \hat{\kappa}_t v_{ii}}}{\Gamma(\hat{\kappa}_t v_{ii}) N_{ii}!}. \tag{37}$$

Further, we see that ξ_{ii} , given data N_{ii} and parameters, has a gamma distribution, i.e.,

$$\xi_{ii} | N_{ii}, \hat{\kappa}_t, \lambda_t \sim \text{Gamma}(\hat{\kappa}_t v_{ii} + N_{ii}, \hat{\kappa}_t v_{ii} + \lambda_t v_{ii}). \tag{38}$$

Given the data augmented likelihood in (37), the estimated parameter $\hat{\kappa}_t$ using (36), and a conjugate gamma prior for λ_t with hyper-parameters $\alpha, \beta > 0$, we can derive the integrated augmented likelihood for terminal node t as follows

$$p_{\text{NB2}}(N_t, \xi_t | \mathbf{X}_t, \mathbf{v}_t, \hat{\kappa}_t) = \frac{\beta^\alpha}{\Gamma(\alpha)} \prod_{i=1}^{n_t} \left[\frac{(\hat{\kappa}_t v_{ii})^{\hat{\kappa}_t v_{ii}} v_{ii}^{N_{ii}}}{\Gamma(\hat{\kappa}_t v_{ii}) N_{ii}!} \xi_{ii}^{\hat{\kappa}_t v_{ii} + N_{ii} - 1} e^{-\xi_{ii} \hat{\kappa}_t v_{ii}} \right] \frac{\Gamma(\sum_{i=1}^{n_t} N_{ii} + \alpha)}{(\sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta)^{\sum_{i=1}^{n_t} N_{ii} + \alpha}}. \tag{39}$$

From the above we see that the posterior distribution of λ_t , given the augmented data (N_t, ξ_t) , is given by

$$\lambda_t | N_t, \xi_t \sim \text{Gamma}\left(\sum_{i=1}^{n_t} N_{ii} + \alpha, \sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta\right).$$

The integrated augmented likelihood for the tree \mathcal{T} is thus given by

$$p_{\text{NB2}}(N, \xi | \mathbf{X}, \mathbf{v}, \hat{\kappa}, \mathcal{T}) = \prod_{t=1}^b p_{\text{NB2}}(N_t, \xi_t | \mathbf{X}_t, \mathbf{v}_t, \hat{\kappa}_t). \tag{40}$$

Now, we discuss the DIC_t for terminal node t of this tree. Similarly, as in the previous subsection, we can easily check that

$$\begin{aligned} \text{DIC}_t &= D(\bar{\lambda}_t) + 2r_{D_t} \\ &= -2 \sum_{i=1}^{n_t} N_{ii} \left(\log(v_{ii}) + \log\left(\sum_{i=1}^{n_t} N_{ii} + \alpha\right) - \log\left(\sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta\right) \right) + 2 \left(\frac{\sum_{i=1}^{n_t} N_{ii} + \alpha}{\sum_{i=1}^{n_t} \xi_{ii} v_{ii} + \beta} \right) \sum_{i=1}^{n_t} \xi_{ii} v_{ii} \\ &\quad + 2 \sum_{i=1}^{n_t} \left(\log(\Gamma(\hat{\kappa}_t v_{ii})) + \log(N_{ii}!) - \hat{\kappa}_t v_{ii} \log(\hat{\kappa}_t v_{ii}) - (\hat{\kappa}_t v_{ii} + N_{ii} - 1) \log(\xi_{ii}) + \xi_{ii} \hat{\kappa}_t v_{ii} \right) \\ &\quad + 2 + 4 \left(\log\left(\sum_{i=1}^{n_t} N_{ii} + \alpha\right) - \psi\left(\sum_{i=1}^{n_t} N_{ii} + \alpha\right) \right) \sum_{i=1}^{n_t} N_{ii}. \end{aligned}$$

For the above two NB models, the DIC of tree \mathcal{T} is obtained by using (10).

With the above formulas derived in the two subsections for NB models, we can use the three-step approach proposed in Section 2.5, together with Algorithm 3, to search for an optimal tree and then obtain predictions for new data.

Algorithm 3 One step of the MCMC algorithm for the NB BCART parameterized by $(\kappa, \lambda, \mathcal{T})$ using data augmentation.

Input: Data (X, v, N) and current values $(\hat{\kappa}^{(m)}, \lambda^{(m)}, \mathcal{T}^{(m)}, \xi^{(m)})$

- 1: Generate a candidate value \mathcal{T}^* with probability distribution $q(\mathcal{T}^{(m)}, \mathcal{T}^*)$
- 2: Estimate $\hat{\kappa}^{(m)}$, using (26) (or (36))
- 3: Sample $\xi^{(m+1)} \sim p(\xi | X, v, N, \hat{\kappa}^{(m+1)}, \lambda^{(m)}, \mathcal{T}^{(m)})$, using (29) (or (38))
- 4: Set the acceptance ratio $\alpha(\mathcal{T}^{(m)}, \mathcal{T}^*) = \min \left\{ \frac{q(\mathcal{T}^*, \mathcal{T}^{(m)}) p_{\text{NB}}(N, \xi^{(m+1)} | X, v, \hat{\kappa}^{(m+1)}, \mathcal{T}^*)}{q(\mathcal{T}^{(m)}, \mathcal{T}^*) p_{\text{NB}}(N, \xi^{(m)} | X, v, \hat{\kappa}^{(m)}, \mathcal{T}^{(m)})} \frac{p(\mathcal{T}^*)}{p(\mathcal{T}^{(m)})}, 1 \right\}$
- 5: Update $\mathcal{T}^{(m+1)} = \mathcal{T}^*$ with probability $\alpha(\mathcal{T}^{(m)}, \mathcal{T}^*)$, otherwise, set $\mathcal{T}^{(m+1)} = \mathcal{T}^{(m)}$
- 6: Sample $\lambda^{(m+1)} \sim \text{Gamma}\left(\sum_{i=1}^{n_t} N_{ii} + \alpha, \sum_{i=1}^{n_t} \xi_{ii}^{(m+1)} v_{ii} + \beta\right)$

Output: New values $(\hat{\kappa}^{(m+1)}, \lambda^{(m+1)}, \mathcal{T}^{(m+1)}, \xi^{(m+1)})$

Remark 7. (a). In step 4 of Algorithm 3, p_{NB} should be understood as either p_{NB1} or p_{NB2} . Similar to Algorithms 1 and 2, the sampling steps in Algorithm 3 should be done when necessary.

(b). It is worth noting that our way of dealing with the parameter κ is different from that in Murray (2021) where a single κ is sampled from a distribution and used for all terminal nodes. It turns out that that way of dealing with κ cannot give us good estimates in our simulation examples, whereas our way of first estimating κ using moment matching method for each node can give good estimates.

(c). There are other ways to parameterize the NB distribution, see, e.g., Zhou et al. (2012). However, it looks that these ways are normally discussed when there is no exposure involved, so we will not cover them here.

3.3. Zero-inflated Poisson models

Insurance claims data normally involves a large volume of zeros. Many policy-holders incur no claims, which does not necessarily mean that they were involved in no accidents, but they are probably less risky. In this section, depending on how the exposure is embedded in the model we discuss two ZIP models to better reflect the excessive zeros, see, e.g., Lee (2021).

3.3.1. Zero-inflated Poisson model 1 (ZIP1)

For terminal node t , we use the following ZIP distribution by embedding the exposure into the Poisson part (see Murray (2021))

$$f_{ZIP1}(m | \mu_t, \lambda_t) = \begin{cases} \frac{1}{1 + \mu_t} + \frac{\mu_t}{1 + \mu_t} f_P(0 | \lambda_t) & m = 0, \\ \frac{\mu_t}{1 + \mu_t} f_P(m | \lambda_t) & m = 1, 2, \dots, \end{cases}$$

$$= \frac{1}{1 + \mu_t} I_{(m=0)} + \frac{\mu_t}{1 + \mu_t} f_P(m | \lambda_t), \quad m = 0, 1, 2, \dots, \tag{41}$$

where $f_P(m | \lambda_t)$ is given as in (16), and $\frac{1}{1 + \mu_t} \in (0, 1)$ is the probability that a zero is due to the point mass component. Note that for computational simplicity we consider a model with two parameters rather than three as in Murray (2021).

Similar to the NB model, a data augmentation scheme is needed for the ZIP model. To this end, we introduce two latent variables $\phi_t = (\phi_{t1}, \phi_{t2}, \dots, \phi_{tn_t}) \in (0, \infty)^{n_t}$ and $\delta_t = (\delta_{t1}, \delta_{t2}, \dots, \delta_{tn_t}) \in \{0, 1\}^{n_t}$, and define the data augmented likelihood for the i -th data instance in terminal node t by

$$f_{ZIP1}(N_{ii}, \delta_{ii}, \phi_{ii} | \mu_t, \lambda_t) = e^{-\phi_{ii}(1 + \mu_t)} \left(\frac{\mu_t (\lambda_t v_{ii})^{N_{ii}}}{N_{ii}!} e^{-\lambda_t v_{ii}} \right)^{\delta_{ii}}, \tag{42}$$

where the support of the function f_{ZIP1} is $(\{0\} \times \{0, 1\} \times (0, \infty)) \cup (\mathbb{N} \times \{1\} \times (0, \infty))$. This means that we impose $\delta_{ii} = 1$ when $N_{ii} \in \mathbb{N}$ (i.e., $N_{ii} \neq 0$). It can be shown that (41) is the marginal distribution of the above augmented distribution; see Murray (2021) for more details. By conditional arguments, we can also check that δ_{ii} , given data $\phi_{ii}, N_{ii} = 0$ and parameters, has a Bernoulli distribution, i.e.,

$$\delta_{ii} | \phi_{ii}, N_{ii} = 0, \mu_t, \lambda_t \sim \text{Bern} \left(\frac{\mu_t e^{-\lambda_t v_{ii}}}{1 + \mu_t e^{-\lambda_t v_{ii}}} \right), \tag{43}$$

and $\delta_{ii} = 1$, given $N_{ii} > 0$. Furthermore, ϕ_{ii} , given data δ_{ii}, N_{ii} and parameters, has an exponential distribution, i.e.,

$$\phi_{ii} | \delta_{ii}, N_{ii}, \mu_t, \lambda_t \sim \text{Exp}(1 + \mu_t). \tag{44}$$

It is noted that the augmented likelihood f_{ZIP1} in (42) can actually be factorized as two gamma-type functions parameterized by μ_t and λ_t respectively. This observation motivates us to assume independent conjugate gamma priors for μ_t and λ_t with hyper-parameters $\alpha_i, \beta_i > 0, i = 1, 2$ (cf. (17)). With these gamma priors, we can derive the integrated augmented likelihood for terminal node t as follows

$$p_{ZIP1}(N_t, \delta_t, \phi_t | X_t, v_t) = \int_0^\infty \int_0^\infty f_{ZIP1}(N_t, \delta_t, \phi_t | \mu_t, \lambda_t) p(\mu_t) p(\lambda_t) d\mu_t d\lambda_t$$

$$= \frac{\beta_1^{\alpha_1}}{\Gamma(\alpha_1)} \frac{\beta_2^{\alpha_2}}{\Gamma(\alpha_2)} \prod_{i=1}^{n_t} \left(e^{-\phi_{ii} v_{ii}^{\delta_{ii}} N_{ii}} (N_{ii}!)^{-\delta_{ii}} \right) \frac{\Gamma(\sum_{i=1}^{n_t} \delta_{ii} + \alpha_1)}{(\sum_{i=1}^{n_t} \phi_{ii} + \beta_1)^{\sum_{i=1}^{n_t} \delta_{ii} + \alpha_1}} \frac{\Gamma(\sum_{i=1}^{n_t} \delta_{ii} N_{ii} + \alpha_2)}{(\sum_{i=1}^{n_t} \delta_{ii} v_{ii} + \beta_2)^{\sum_{i=1}^{n_t} \delta_{ii} N_{ii} + \alpha_2}}. \tag{45}$$

Moreover, from the above we see that the posterior distributions of μ_t, λ_t given the augmented data (N_t, δ_t, ϕ_t) are given by

$$\mu_t | N_t, \delta_t, \phi_t \sim \text{Gamma} \left(\sum_{i=1}^{n_t} \delta_{ii} + \alpha_1, \sum_{i=1}^{n_t} \phi_{ii} + \beta_1 \right),$$

$$\lambda_t | N_t, \delta_t, \phi_t \sim \text{Gamma} \left(\sum_{i=1}^{n_t} \delta_{ii} N_{ii} + \alpha_2, \sum_{i=1}^{n_t} \delta_{ii} v_{ii} + \beta_2 \right).$$

The integrated augmented likelihood for the tree \mathcal{T} is thus given by

$$p_{ZIP1}(N, \delta, \phi | X, v, \mathcal{T}) = \prod_{t=1}^b p_{ZIP1}(N_t, \delta_t, \phi_t | X_t, v_t). \tag{46}$$

Now, we discuss the DIC for this tree which can be derived as a special case of (11) with $\theta_t = (\mu_t, \lambda_t)$. To this end, we first focus on the DIC_t of terminal node t . It follows that

$$D(\bar{\mu}_t, \bar{\lambda}_t) = -2 \log f_{ZIP1}(N_t | \bar{\mu}_t, \bar{\lambda}_t) = -2 \sum_{i=1}^{n_t} \log \left(\frac{1}{1 + \bar{\mu}_t} I_{(N_{ii}=0)} + \frac{\bar{\mu}_t}{1 + \bar{\mu}_t} \frac{(\bar{\lambda}_t v_{ii})^{N_{ii}}}{N_{ii}!} e^{-\bar{\lambda}_t v_{ii}} \right), \tag{47}$$

where

$$\bar{\mu}_t = \frac{\sum_{i=1}^{n_t} \delta_{ii} + \alpha_1}{\sum_{i=1}^{n_t} \phi_{ii} + \beta_1}, \quad \bar{\lambda}_t = \frac{\sum_{i=1}^{n_t} \delta_{ii} N_{ii} + \alpha_2}{\sum_{i=1}^{n_t} \delta_{ii} v_{ii} + \beta_2}.$$

Next, since

$$\log f_{ZIP1}(\mathbf{N}_t, \boldsymbol{\delta}_t, \boldsymbol{\phi}_t \mid \mu_t, \lambda_t) = \sum_{i=1}^{n_t} (-\phi_{ti}(1 + \mu_t) + \delta_{ti} \log(\mu_t) + \delta_{ti} N_{ti} \log(\lambda_t v_{ti}) - \delta_{ti} \lambda_t v_{ti} - \delta_{ti} \log(N_{ti}!)),$$

we can derive that

$$\begin{aligned} q_{Dt} &= -2E_{\text{post}}(\log f_{ZIP1}(\mathbf{N}_t, \boldsymbol{\delta}_t, \boldsymbol{\phi}_t \mid \mu_t, \lambda_t)) + 2 \log f_{ZIP1}(\mathbf{N}_t, \boldsymbol{\delta}_t, \boldsymbol{\phi}_t \mid \bar{\mu}_t, \bar{\lambda}_t) \\ &= 2 \left(\log \left(\sum_{i=1}^{n_t} \delta_{ti} + \alpha_1 \right) - \psi \left(\sum_{i=1}^{n_t} \delta_{ti} + \alpha_1 \right) \right) \sum_{i=1}^{n_t} \delta_{ti} + 2 \left(\log \left(\sum_{i=1}^{n_t} \delta_{ti} N_{ti} + \alpha_2 \right) - \psi \left(\sum_{i=1}^{n_t} \delta_{ti} N_{ti} + \alpha_2 \right) \right) \sum_{i=1}^{n_t} \delta_{ti} N_{ti}. \end{aligned} \tag{48}$$

Therefore, DIC_t can be obtained from (47) and (48) as

$$DIC_t = D(\bar{\mu}_t, \bar{\lambda}_t) + 2q_{Dt}.$$

3.3.2. Zero-inflated Poisson model 2 (ZIP2)

For terminal node t , we use the following ZIP distribution by embedding the exposure into the zero mass part (see Lee (2021))

$$f_{ZIP2}(m \mid \mu_t, \lambda_t) = \begin{cases} \frac{1}{1 + \mu_t v_{ti}} + \frac{\mu_t v_{ti}}{1 + \mu_t v_{ti}} e^{-\lambda_t} & m = 0, \\ \frac{\mu_t v_{ti}}{1 + \mu_t v_{ti}} \frac{\lambda_t^{N_{ti}}}{N_{ti}!} e^{-\lambda_t} & m = 1, 2, \dots, \end{cases} \tag{49}$$

where $\frac{1}{1 + \mu_t v_{ti}} \in (0, 1)$ is the probability that a zero is due to the point mass component. This formulation stems from an intuitive inverse relationship between the exposure and the probability of zero mass. This way of embedding exposure has been justified to be more effective in Lee (2021).

Similar to before, we introduce two latent variables $\boldsymbol{\phi}_t = (\phi_{t1}, \phi_{t2}, \dots, \phi_{tn_t}) \in (0, \infty)^{n_t}$ and $\boldsymbol{\delta}_t = (\delta_{t1}, \delta_{t2}, \dots, \delta_{tn_t}) \in \{0, 1\}^{n_t}$, and define the data augmented likelihood for the i -th data instance in terminal node t as

$$f_{ZIP2}(N_{ti}, \delta_{ti}, \phi_{ti} \mid \mu_t, \lambda_t) = e^{-\phi_{ti}(1 + \mu_t v_{ti})} \left(\frac{\mu_t v_{ti} \lambda_t^{N_{ti}}}{N_{ti}!} e^{-\lambda_t} \right)^{\delta_{ti}}, \tag{50}$$

where the support of the function f_{ZIP2} is $(\{0\} \times \{0, 1\} \times (0, \infty)) \cup (\mathbb{N} \times \{1\} \times (0, \infty))$. By conditional arguments, we can also check that δ_{ti} , given data $\phi_{ti}, N_{ti} = 0$ and parameters, has a Bernoulli distribution, i.e.,

$$\delta_{ti} \mid \phi_{ti}, N_{ti} = 0, \mu_t, \lambda_t \sim \text{Bern} \left(\frac{\mu_t v_{ti} e^{-\lambda_t}}{1 + \mu_t v_{ti} e^{-\lambda_t}} \right), \tag{51}$$

and $\delta_{ti} = 1$, given $N_{ti} > 0$. Furthermore, ϕ_{ti} , given data δ_{ti}, N_{ti} and parameters, has an exponential distribution, i.e.,

$$\phi_{ti} \mid \delta_{ti}, N_{ti}, \mu_t, \lambda_t \sim \text{Exp}(1 + \mu_t v_{ti}). \tag{52}$$

As previously, we assume independent conjugate gamma priors for μ_t and λ_t with hyper-parameters $\alpha_i, \beta_i > 0, i = 1, 2$. Given the data augmented likelihood in (50) and the above gamma priors, we can obtain the integrated augmented likelihood for terminal node t as follows

$$p_{ZIP2}(\mathbf{N}_t, \boldsymbol{\delta}_t, \boldsymbol{\phi}_t \mid \mathbf{X}_t, \mathbf{v}_t) = \frac{\beta_1^{\alpha_1}}{\Gamma(\alpha_1)} \frac{\beta_2^{\alpha_2}}{\Gamma(\alpha_2)} \prod_{i=1}^{n_t} \left(e^{-\phi_{ti}} \left(\frac{v_{ti}}{N_{ti}!} \right)^{\delta_{ti}} \right) \frac{\Gamma(\sum_{i=1}^{n_t} \delta_{ti} + \alpha_1)}{(\sum_{i=1}^{n_t} \phi_{ti} v_{ti} + \beta_1)^{\sum_{i=1}^{n_t} \delta_{ti} + \alpha_1}} \frac{\Gamma(\sum_{i=1}^{n_t} \delta_{ti} N_{ti} + \alpha_2)}{(\sum_{i=1}^{n_t} \delta_{ti} + \beta_2)^{\sum_{i=1}^{n_t} \delta_{ti} N_{ti} + \alpha_2}}. \tag{53}$$

Moreover, from the above we see that the posterior distributions of μ_t, λ_t given the augmented data $(\mathbf{N}_t, \boldsymbol{\delta}_t, \boldsymbol{\phi}_t)$ are given by

$$\begin{aligned} \mu_t \mid \mathbf{N}_t, \boldsymbol{\delta}_t, \boldsymbol{\phi}_t &\sim \text{Gamma} \left(\sum_{i=1}^{n_t} \delta_{ti} + \alpha_1, \sum_{i=1}^{n_t} \phi_{ti} v_{ti} + \beta_1 \right), \\ \lambda_t \mid \mathbf{N}_t, \boldsymbol{\delta}_t, \boldsymbol{\phi}_t &\sim \text{Gamma} \left(\sum_{i=1}^{n_t} \delta_{ti} N_{ti} + \alpha_2, \sum_{i=1}^{n_t} \delta_{ti} + \beta_2 \right). \end{aligned}$$

The integrated augmented likelihood for the tree \mathcal{T} is thus given by

$$p_{ZIP2}(\mathbf{N}, \boldsymbol{\delta}, \boldsymbol{\phi} \mid \mathbf{X}, \mathbf{v}, \mathcal{T}) = \prod_{t=1}^b p_{ZIP2}(\mathbf{N}_t, \boldsymbol{\delta}_t, \boldsymbol{\phi}_t \mid \mathbf{X}_t, \mathbf{v}_t, \mathcal{T}). \tag{54}$$

Now, we discuss the DIC_t of terminal node t . It follows that

$$D(\bar{\mu}_t, \bar{\lambda}_t) = -2 \log f_{ZIP2}(\mathbf{N}_t \mid \bar{\mu}_t, \bar{\lambda}_t) = -2 \sum_{i=1}^{n_t} \log \left(\frac{1}{1 + \bar{\mu}_t v_{ti}} I_{(N_{ti}=0)} + \frac{\bar{\mu}_t v_{ti}}{1 + \bar{\mu}_t v_{ti}} \frac{\bar{\lambda}_t^{N_{ti}}}{N_{ti}!} e^{-\bar{\lambda}_t} \right), \tag{55}$$

where

$$\bar{\mu}_t = \frac{\sum_{i=1}^{n_t} \delta_{ti} + \alpha_1}{\sum_{i=1}^{n_t} \phi_{ti} v_{ti} + \beta_1}, \quad \bar{\lambda}_t = \frac{\sum_{i=1}^{n_t} \delta_{ti} N_{ti} + \alpha_2}{\sum_{i=1}^{n_t} \delta_{ti} + \beta_2}.$$

Next, since

$$\log f_{ZIP2}(N_t, \delta_t, \phi_t | \mu_t, \lambda_t) = \sum_{i=1}^{n_t} (-\phi_{ti}(1 + \mu_t v_{ti}) + \delta_{ti} \log(\mu_t v_{ti}) + \delta_{ti} N_{ti} \log(\lambda_t) - \delta_{ti} \lambda_t - \delta_{ti} \log(N_{ti}!)),$$

we can derive the same expression for q_{Dt} as in (48). Therefore, we obtain from (55) and (48) that

$$\begin{aligned} DIC_t &= -2 \sum_{i=1}^{n_t} \log \left(\frac{1}{1 + \bar{\mu}_t v_{ti}} I_{(N_{ti}=0)} + \frac{\bar{\mu}_t v_{ti}}{1 + \bar{\mu}_t v_{ti}} \frac{\bar{\lambda}_t^{N_{ti}}}{N_{ti}!} e^{-\bar{\lambda}_t} \right) \\ &+ 4 \left(\log \left(\sum_{i=1}^{n_t} \delta_{ti} + \alpha_1 \right) - \psi \left(\sum_{i=1}^{n_t} \delta_{ti} + \alpha_1 \right) \right) \sum_{i=1}^{n_t} \delta_{ti} \\ &+ 4 \left(\log \left(\sum_{i=1}^{n_t} \delta_{ti} N_{ti} + \alpha_2 \right) - \psi \left(\sum_{i=1}^{n_t} \delta_{ti} N_{ti} + \alpha_2 \right) \right) \sum_{i=1}^{n_t} \delta_{ti} N_{ti}. \end{aligned}$$

For the above two ZIP models, the DIC of tree \mathcal{T} is obtained using (10).

With the formulas derived in the above two subsections for ZIP models, we can use the three-step approach proposed in Section 2.5, together with Algorithm 2, to search for an optimal tree and then obtain predictions for new data.

Remark 8. There are other ways to deal with the data augmentation for ZIP models; see, e.g., Tanner and Wong (1987); Rodrigues (2003); Diebolt and Robert (1994) where only one latent variable is introduced. The models discussed therein with one latent variable should work more efficiently, but in their constructions no exposure is considered. Since involvement of exposure is one of the key features of insurance claims frequency analysis, we had to introduce two latent variables for data augmentation to facilitate calculations.

4. Simulation and real data analysis

In this section, we illustrate the efficiency of the BCART models introduced in Section 3 by using simulated data and a real insurance claims dataset. In the sequel, we use the abbreviation P-CART to denote CART for the Poisson model, and other abbreviations can be similarly understood (e.g., NB1-BCART denotes the BCART for NB model 1).

4.1. Performance measures

We first introduce some performance measures that will be used for prediction comparisons. Suppose we have obtained a tree with b terminal nodes and the corresponding parameter estimates for θ which we will use to obtain the prediction \hat{N}_i given x_i, v_i for a test data set (X, v, N) with m observations. The number of test data in the terminal node t is denoted by $m_t, t = 1, \dots, b$. The performance measures used are as follows:

M1: The residual sum of squares (RSS) is given by

$$RSS(N) = \sum_{i=1}^m (N_i - \hat{N}_i)^2.$$

This measure is commonly used for Gaussian-distributed data, but here we also use it for non-Gaussian data for comparison.

M2: RSS based on a sub-portfolio (i.e., those instances in the same terminal node) level is given by

$$RSS(N/v) = \sum_{t=1}^b \left(\frac{\sum_{i=1}^{m_t} N_{ti}}{\sum_{i=1}^{m_t} v_{ti}} - \hat{y}_t \right)^2,$$

where \hat{y}_t is the estimated frequency for the terminal node t , which is estimated by (14) assuming unit exposure. More specifically, $\hat{y}_t = \bar{\lambda}_t$ for Poisson and NB models, and $\hat{y}_t = \bar{\mu}_t \bar{\lambda}_t (1 + \bar{\mu}_t)^{-1}$ for ZIP models. This measure is preferred here as it takes account of accuracy on a (sub-)portfolio level (i.e., balance property) other than an individual level. We refer to Denuit et al. (2021); Wüthrich (2020, 2022) for more details and discussions of the balance property that is required for insurance pricing.

M3: Negative log-likelihood (NLL): This is calculated by using the assumed response distribution in the terminal node with the estimated parameters. It represents the ex-ante belief of the underlying distribution of the data, and is thus a good measure for model comparison, see, e.g., Lee (2021).

M4: Discrepancy statistic (DS) (cf. Naya et al. (2008)), is defined as a weighted version of $RSS(N/v)$, given by

$$DS(N/v) = \sum_{t=1}^b \frac{1}{\hat{\sigma}_t^2} \left(\frac{\sum_{i=1}^{m_t} N_{ti}}{\sum_{i=1}^{m_t} v_{ti}} - \hat{y}_t \right)^2,$$

where \hat{y}_t is the same as in M2, and $\hat{\sigma}_t^2$ is the estimated variance of frequency of terminal node t . More specifically, $\hat{\sigma}_t^2 = \bar{\lambda}_t$ for the Poisson model, $\hat{\sigma}_t^2 = \bar{\lambda}_t (1 + \bar{\lambda}_t / \bar{\kappa}_t)$ for NB models, and $\hat{\sigma}_t^2 = \bar{\mu}_t \bar{\lambda}_t (1 + \bar{\mu}_t + \bar{\lambda}_t) (1 + \bar{\mu}_t)^{-2}$ for ZIP models.

M5: Lift: Model lift indicates the ability to differentiate between low and high claims frequency policy-holders. Sometimes it is called the “economic value” of the model. A higher lift illustrates that the model is more capable of separating the extreme values from the average. We refer to Henckaerts et al. (2021); Lee (2020, 2021) and references therein for further discussion on lift. We propose a way to calculate lift for the tree model in the following steps.

Step 1: Retrieve the predicted frequencies for terminal nodes, $\hat{y}_t, t = 1, \dots, b$, for the optimal tree obtained from the training procedure.

Step 2: Set $\hat{y}_{\min} = \min_{t=1}^b \hat{y}_t$ and $\hat{y}_{\max} = \max_{t=1}^b \hat{y}_t$, which identify the least and most risky groups of policy-holders, respectively.

Table 2

Total count each variable used amongst all accepted trees from the P-BCART MCMC algorithms (after burn-in period; equal probabilities for tree moves; one run with 3 restarts).

γ	ρ	# terminal nodes	# accepted trees	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8
0.50	20	2	342	197	156	1	0	2	5	4	4
0.95	17	3	460	408	381	1	4	8	6	8	5
0.99	15	4	800	1261	1239	12	17	30	25	31	20
0.99	12	5	652	1157	1126	9	7	18	15	16	20
0.99	10	6	305	710	680	13	4	18	25	30	12
0.99	6	7	318	825	809	3	8	15	23	14	9
0.99	5	8	210	681	647	2	10	7	13	18	5

Step 3: Use test data in the least and most risky groups/nodes to obtain their total sum of exposures, say v_{\min} and v_{\max} .

Step 4: If $v_{\min} \leq v_{\max}$, then sort the data using exposures in descending order in the most risky group. Calculate the cumulative sums of the sorted exposures until the one equal or greater than v_{\min} is achieved and then calculate the corresponding empirical frequency (i.e., ratio of sum of claim numbers and sum of exposures) of these first data involved, say $\lambda_{\max}^{(e)}$. The lift is defined as $L = \lambda_{\max}^{(e)} / \lambda_{\min}^{(e)}$, where

$\lambda_{\min}^{(e)}$ is the empirical frequency of the least risky group.

[Similarly, If $v_{\min} > v_{\max}$, then sort the data using exposures in ascending order in the least risky group. Calculate the cumulative sums of the sorted exposures until the one equal or greater than v_{\max} is achieved and then calculate the corresponding empirical frequency of these first data involved, say $\lambda_{\min}^{(e)}$. The lift is defined as $L = \lambda_{\max}^{(e)} / \lambda_{\min}^{(e)}$, where $\lambda_{\max}^{(e)}$ is the empirical frequency of the most risky group.]

We remark that more performance measures and diagnostic approaches can be introduced following ideas in, e.g., Lee (2021, 2020); Henckaerts et al. (2021). However, this is not the main focus of the present paper, so these are explored elsewhere.

4.2. Simulation examples

We will discuss three simulation examples, namely, Scenarios 1–3 below. Scenario 1 aims to illustrate that BCART can do really well for the chessboard data similar to Fig. 1 for which CART cannot reasonably do anything. In addition, from this simulation study we also see that BCART can do well with variable selection. In Scenario 2, we shall examine how different BCART models can capture the data over-dispersion. In Scenario 3, we illustrate the effectiveness of ZIP-BCART models for data with exposures.

4.2.1. Scenario 1: Poisson data with noise variables

We simulate a data set $\{(x_i, v_i, N_i)\}_{i=1}^n$ with $n = 5,000$ independent observations. Here $v_i \sim U(0, 1)$, $x_i = (x_{i1}, \dots, x_{i8})$, with independent components $x_{i1} \sim U\{-3, -2, -1, 1, 2, 3\}$, $x_{i2} \sim N(0, 1)$, $x_{ik} \sim U(-1, 1)$ for $k = 3, 4$, $x_{ik} \sim N(0, 1)$ for $k = 5, 6$, and $x_{ik} \sim U\{-3, -2, -1, 1, 2, 3\}$ for $k = 7, 8$. Moreover, $N_i \sim \text{Poi}(\lambda(x_{i1}, x_{i2}) v_i)$, where

$$\lambda(x_1, x_2) = \begin{cases} 1 & \text{if } x_1 x_2 \leq 0, \\ 7 & \text{if } x_1 x_2 > 0. \end{cases}$$

Obviously, the designed noise variables $x_{ik}, k = 3, \dots, 8$ are all independent of the response N . We use P-BCART and P-CART for the above simulated data, where $x_{ik}, k = 1, 7, 8$ are treated as categorical. We have included both categorical and continuous variables as noise variables and as significant variables, which is a bit more general than the data shown in Fig. 1. Note that the same conclusion can be drawn for numeric $x_{ik}, k = 1, 7, 8$, but to better illustrate the effectiveness of the P-BCART we choose to make them as characters (to increase the splitting possibilities of these variables).

We first apply P-CART as implemented in the R package `rpart` (Therneau and Atkinson, 2019). It is not surprising that P-CART is not able to give us any reasonable tree that can characterize the data, due to its greedy search nature. The smallest tree (except the one with only a root node) that P-CART generated has 25 terminal nodes and the tree found by using cross-validation has 31 terminal nodes. Obviously, both of them are much more complicated than the real model. Furthermore, in these two trees all the noise variables are used, which indicates that P-CART is sensitive to noise.

Now we discuss the P-BCART applied to the data focusing preliminary on the effect of noise variables to the model. We simply set equal probabilities, i.e., $P(\text{Grow}) = 0.2$, $P(\text{Prune}) = 0.2$, $P(\text{Change1}) = 0.2$, $P(\text{Change2}) = 0.2$ and $P(\text{Swap}) = 0.2$, for the tree proposals. For the gamma prior of the Poisson intensities λ_i we use $\alpha = 3.2096$ and $\beta = 0.8$ which are selected by keeping the relationship $\alpha/\beta = \sum_{i=1}^n N_i / \sum_{i=1}^n v_i$. It is worth mentioning that the performance of the algorithm does not change much when choosing different pairs of (α, β) while keeping their ratio. We also observe the same in other simulation examples, so in the following we will not dwell on their selection.

In Table 2 we list the tuned hyper-parameters γ, ρ in the first two columns for which the MCMC algorithms will converge to a region of trees with a certain number of terminal nodes listed (see Step 1 of Table 1). For each fixed hyper-parameter γ and ρ , we run 10000 iterations in the MCMC algorithm and take results after an initial burn-in period of 2000 iterations, after which the posterior probabilities of the tree structures have been settled for some time. This procedure is done with 3 restarts. The fourth column gives the total number of accepted trees after the burn-in period in the MCMC algorithms. The last columns of Table 2 include the total number of times each variable is used in the accepted trees. We see from these columns that all noise variables have a very low selection rate, and as expected, the significant variables x_1, x_2 are dominating. Besides, at a first glance, it is inferred that the noise variables x_3 and x_4 have a much lower selection rate than the other noise variables which is just because x_3 and x_4 are simulated using a distribution completely different from those of the significant variables. However, when the experiment is run 10 times, we find that the average selection rates of all noise variables are almost the same independent of their distributions (see Table 3), which is consistent with the expectation.

In Fig. 2, we illustrate this procedure for $j = 4$ (the same as that summarized in the third row of Table 2), with plots of the number of terminal nodes, the integrated likelihood $p_p(N|X, v, T)$ and the data likelihood $p_p(N|X, v, \bar{\lambda}, T)$ of the accepted trees. The observations are in line with

Table 3

Average frequency each variable used in all accepted trees from the P-BCART MCMC algorithms (after burn-in period; equal probabilities for tree moves; ten runs with 3 restarts).

# terminal nodes	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8
2	183	140	1	1	1	3	2	1
3	422	405	2	3	4	3	3	2
4	1242	1201	11	13	15	13	14	12
5	1207	1162	12	14	12	13	11	14
6	821	828	9	8	10	12	11	8
7	998	976	8	9	10	12	10	8
8	847	795	7	9	9	11	10	8

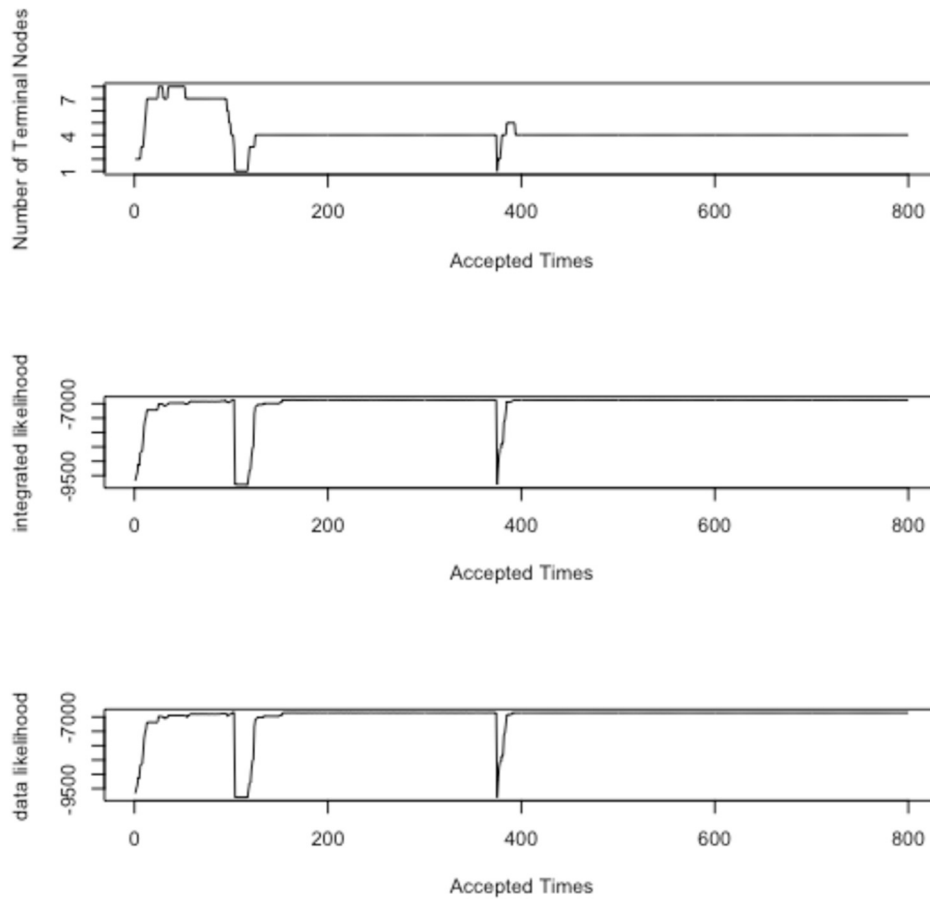


Fig. 2. Trace plots from MCMC with 3 restarts ($\gamma = 0.99, \rho = 15$).

those in Chipman et al. (1998); we see from the likelihood plots that the convergence of MCMC can be obtained relatively quickly. Interestingly, the optimal tree is not found in the first round of MCMC which got stuck in a local mode, but the restarts helped where in the second and the third rounds optimal trees can be found. Moreover, we see that there is no big difference shown in the plots of the integrated likelihood and the data likelihood.

Following Step 2 of Table 1, for each $j = 2, \dots, 8$, we select the optimal tree with maximum data likelihood $p_D(\mathbf{N}|\mathbf{X}, \mathbf{v}, \bar{\lambda}, \mathcal{T})$ from the convergence region. The variables used in these optimal trees are listed in Table 4, where we can see that none of these trees involves any of the noise variables. The values for the effective number of parameters p_D reflect the number of parameters in the tree if a flat prior for λ_i is used. Furthermore, we list the DIC for these trees in the last column of Table 4. Following Step 3 of Table 1 we conclude that the selected optimal tree is the one with 4 terminal nodes which is illustrated in Fig. 3. We see that this tree is close to a true optimal one with the almost correct topology and accurate parameter estimates.

Using equal probabilities for the proposed tree moves, the above example provides detailed information about how to implement the three-step tree selection procedure in practice and illustrates the effectiveness of the method. Next, we investigate which type of step (particularly, the Change and Swap moves) contributes more to the computational efficiency. To this end, we shall vary the probabilities of the Change and Swap moves, keeping the same probabilities for Grow and Prune moves at 0.2. Different experiments can be designed as in Table 5.

We fix $\gamma = 0.99$ and $\rho = 15$, as for Fig. 2. For each of the experiments E1–E4, we run the P-BCART MCMC algorithm 10 times and for each run we record the iteration time until an “optimal” tree is found. The average iteration time with the standard deviation (s.d.) of the 10 runs and the average acceptance rates of moves are shown in Table 6. The figures in the second row indicate that the experiment E4 is faster in finding an “optimal” tree than E1–E3 when at least one of the Change moves or/and the Swap move is removed. In particular, the comparison between E1 and E2 confirms

Table 4

Number of times each variable used in each chosen optimal tree and the corresponding p_D and DIC (after burn-in period; equal probabilities for tree moves; one run with 3 restarts). Bold font indicates DIC selected model.

# terminal nodes	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	p_D	DIC
2	1	0	0	0	0	0	0	0	2.00	14221
3	1	1	0	0	0	0	0	0	2.95	14076
4	1	2	0	0	0	0	0	0	3.97	13526
5	2	2	0	0	0	0	0	0	4.97	13570
6	3	2	0	0	0	0	0	0	5.93	13629
7	3	3	0	0	0	0	0	0	6.91	13678
8	4	3	0	0	0	0	0	0	7.95	13683

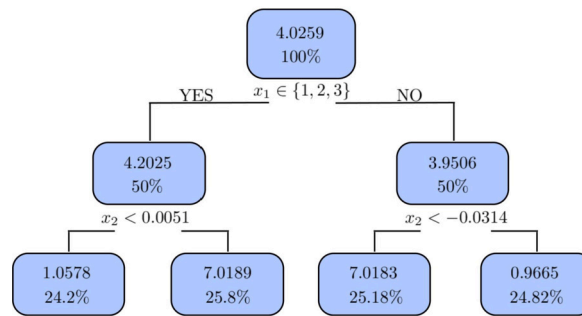


Fig. 3. Optimal P-BCART. Numbers at each node give the estimated value for the frequency parameter λ , and the percentage of observations.

Table 5

Four different experiments (E1–E4) for given probabilities of tree moves. In each case probability of Grow and Prune is fixed at 0.2.

	Change1	Change2	Swap
E1	0	0.6	0
E2	0	0.3	0.3
E3	0.3	0	0.3
E4	0.2	0.2	0.2

Table 6

Average iteration times to obtain an “optimal” tree (4 terminal nodes) and accepted move rates from the P-BCART MCMC algorithms (after burn-in period; ten runs with 3 restarts). Experiments E1–E4 are described in Table 5.

	E1	E2	E3	E4
Average iteration times (s.d.)	3388 (168)	2710 (187)	2984 (177)	2018 (161)
Acceptance rate of all moves	3.10%	3.23%	3.14%	3.87%
Acceptance rate of Grow	1.50%	1.36%	0.90%	0.65%
Acceptance rate of Prune	1.43%	1.20%	0.54%	0.30%
Acceptance rate of Change1	-	-	6.09%	8.19%
Acceptance rate of Change2	4.17%	4.87%	-	5.66%
Acceptance rate of Swap	-	4.01%	3.49%	4.52%

the essence of the Swap move, as illustrated also in Chipman et al. (1998). Moreover, the acceptance rate of all moves is a weighted average of acceptance rates of all individual moves, and we observe that the acceptance rates of the Change and Swap moves (in particular, the Change1 move) are significantly greater than Grow and Prune moves, which also confirms the significance of the Change and Swap moves (especially, the Change1 move).

We also ran several other similar but more complex simulation examples to check the performance of P-BCART, NB-BCART and ZIP-BCART models. Our conclusions from these simulations are: 1) BCART models can retrieve the tree structure (including both topology and parameters) as that used to simulate the data, 2) BCART models are able to avoid choosing noise variables regardless of their distributions, and 3) the Change and Swap moves have significant impacts on the BCART models and it is beneficial to include two types of the Change move.

4.2.2. Scenario 2: ZIP data with varying probability of zero mass component

We simulate a data set $\{(x_i, v_i, N_i)\}_{i=1}^n$ with $n = 5,000$ independent observations. Here $x_i = (x_{i1}, x_{i2})$, with $x_{ik} \sim N(0, 1)$ for $k = 1, 2$. We assume $v_i \equiv 1$ for simplicity, since it is not a key feature in this Scenario. Moreover, $N_i \sim \text{ZIP}(p_0, \lambda(x_{i1}, x_{i2}))$, where

$$\lambda(x_1, x_2) = \begin{cases} 7 & \text{if } x_1 x_2 \leq 0, \\ 1 & \text{if } x_1 x_2 > 0, \end{cases}$$

Table 7
Hyper-parameters, p_D (or q_D, r_D) and DIC on training data ($p_0 = 0.05$). Bold font indicates DIC selected model.

Model	γ	ρ	p_D (or q_D, r_D)	DIC
ZIP-BCART (2)	0.50	20	4.00	11451
ZIP-BCART (3)	0.99	20	5.94	11405
ZIP-BCART (4)	0.99	15	7.95	11322
ZIP-BCART (5)	0.99	5	9.86	11364
P-BCART (2)	0.50	20	2.00	11369
P-BCART (3)	0.99	20	2.99	11337
P-BCART (4)	0.99	10	3.99	11262
P-BCART (5)	0.99	5	4.91	11299
NB-BCART (2)	0.50	30	4.00	11317
NB-BCART (3)	0.99	25	5.99	11273
NB-BCART (4)	0.99	20	7.99	11192
NB-BCART (5)	0.99	5	9.90	11237

Table 8
Model performance on test data ($p_0 = 0.05$) with bold entries determined by DIC (see Table 7).

Model	RSS(N)	RSS(N/ν)	NLL	DS(N/ν)	Lift
ZIP-BCART (2)	2013	0.00222	1975	0.000185	1.22
ZIP-BCART (3)	1986	0.00208	1953	0.000169	2.67
ZIP-BCART (4)	1923	0.00162	1890	0.000116	6.34
ZIP-BCART (5)	1909	0.00182	1863	0.000130	6.56
P-BCART (2)	1758	0.00175	1702	0.000138	1.40
P-BCART (3)	1732	0.00160	1673	0.000123	3.21
P-BCART (4)	1681	0.00108	1612	0.000072	6.62
P-BCART (5)	1662	0.00126	1594	0.000092	6.75
NB-BCART (2)	1683	0.00145	1647	0.000101	1.58
NB-BCART (3)	1661	0.00131	1616	0.000092	3.53
NB-BCART (4)	1609	0.00070	1536	0.000056	6.95
NB-BCART (5)	1589	0.00097	1502	0.000074	6.97

and $p_0 \in (0, 1)$ is the probability of a zero due to the point mass component, for which the value is to be specified. The data is split into two subsets: a training set with $n - m = 4,000$ observations and a test set with $m = 1,000$ observations.

For this Scenario, we aim to examine how the P-BCART, NB-BCART and ZIP-BCART will perform when p_0 is varied. Note that since $\nu_i \equiv 1$, NB1 and NB2 (ZIP1 and ZIP2) will be essentially the same. Intuition tells us that when p_0 is small NB-BCART should be good enough to capture the over-dispersion introduced by a small proportion of zeros, but when p_0 becomes large ZIP-BCART should perform better for the highly over-dispersed data. This intuition will be confirmed by this study. For simplicity, we shall present two results, one with $p_0 = 0.05$ and the other with $p_0 = 0.95$.

We first discuss the simulation with a small probability of zero mass (i.e., $p_0 = 0.05$). In Table 7 we present the hyper-parameters γ, ρ used to obtain MCMC convergence to the region of trees with a certain number of terminal nodes (indicated after the abbreviation of models, e.g., the 2 in ZIP-BCART (2)). The last two columns give the effective number of parameters and DIC of the optimal trees for each model, respectively. We can conclude from the DIC that by using Step 3 in Table 1 we can select the optimal tree with the true 4 terminal nodes for either ZIP-BCART, P-BCART or NB-BCART, and among those, the NB-BCART (with DIC = 11192) is the best one. This looks a bit surprising at a first glance because our data are simulated from a ZIP model. We suspect that the reason for this may be two-fold: First, the NB is enough to capture the small over-dispersion. Second, we have used data-augmentation in the algorithms and thus it is understandable that the NB-BCART with 1 latent variable (see Section 3.2) could achieve better performance than the “real” ZIP-BCART with 2 latent variables (see Section 3.3). Moreover, we see that even the P-BCART performs better than the ZIP-BCART, for similar reasons.

Now, let us look at the performance of these models on test data in Table 8. First, we see that for each type of model, ZIP, Poisson and NB, the optimal tree with 4 terminal nodes achieves best RSS(N/ν) (0.00162, 0.00108 and 0.00070, respectively) and DS(N/ν) (0.000116, 0.000072 and 0.000056, respectively) on test data, which is not surprising as those models retrieve the almost true tree structures. Second, we see from RSS(N) that for each type of model, the performance becomes better as the number of terminal nodes that we want increases, however, the amount of decrement becomes smaller after the optimal trees with 4 terminal nodes have been obtained. We observe the same for negative log-likelihood and lift. It is worth noting that when calculating and comparing lift for different trees, instead of simply following the four steps in M5, in Step 4 we first choose the minimum total sum of exposures among the least and most risky groups in all the trees to be compared and then calculate other values accordingly using this minimum total sum of exposures as the basis. Third, we see that among these three trees with 4 terminal nodes, the one obtained from NB-BCART gives the best performance on test data based on all these performance measures, which is consistent with the conclusion from training data.

Next, we consider the simulation with a large probability of zero mass (i.e., $p_0 = 0.95$). The results are displayed in Tables 9 and 10. Similar discussions can be done for this case. In particular, we find that the performance order based on DIC is ZIP-BCART > NB-BCART > P-BCART, which is also consistent with their performance on test data.

We also ran several other similar simulation examples to check the performance of P-BCART, NB-BCART and ZIP-BCART with different values for p_0 . Our conclusion from these simulations is that when the proportion of zeros in the data is small (reflected by small p_0) then the NB-BCART or

Table 9
Hyper-parameters, p_D (or q_D, r_D) and DIC on training data ($p_0 = 0.95$). Bold font indicates DIC selected model.

Model	γ	ρ	p_D (or q_D, r_D)	DIC
ZIP-BCART (2)	0.50	10	3.99	3483
ZIP-BCART (3)	0.99	10	5.99	3452
ZIP-BCART (4)	0.99	8	7.95	3375
ZIP-BCART (5)	0.99	3	9.93	3396
P-BCART (2)	0.50	10	1.98	3892
P-BCART (3)	0.99	10	2.96	3863
P-BCART (4)	0.99	5	3.91	3801
P-BCART (5)	0.99	2	4.90	3827
NB-BCART (2)	0.50	20	3.99	3726
NB-BCART (3)	0.99	20	5.97	3699
NB-BCART (4)	0.99	10	7.92	3632
NB-BCART (5)	0.99	8	9.89	3667

Table 10
Model performance on test data ($p_0 = 0.95$) with bold entries determined by DIC (see Table 9).

Model	RSS(N)	RSS(N/v)	NLL	DS(N/v)	Lift
ZIP-BCART (2)	721	0.00755	699	0.00721	1.25
ZIP-BCART (3)	715	0.00700	690	0.00698	1.92
ZIP-BCART (4)	682	0.00571	657	0.00619	2.86
ZIP-BCART (5)	675	0.00613	649	0.00646	3.13
P-BCART (2)	782	0.00967	754	0.00802	1.15
P-BCART (3)	773	0.00891	746	0.00786	1.50
P-BCART (4)	750	0.00723	719	0.00712	2.40
P-BCART (5)	741	0.00792	705	0.00739	2.72
NB-BCART (2)	775	0.00893	740	0.00775	1.19
NB-BCART (3)	768	0.00810	731	0.00740	1.72
NB-BCART (4)	735	0.00647	701	0.00667	2.60
NB-BCART (5)	730	0.00703	693	0.00689	2.90

P-BCART performs better than ZIP-BCART, whereas when the proportion of zeros in the data is large then the ZIP-BCART is preferred to NB-BCART and P-BCART. This finding is consistent with the real insurance data discussed below.

4.2.3. Scenario 3: different ways to incorporate exposure in ZIP models

The purpose of Scenario 3 is to compare two different ways of dealing with exposure, namely, ZIP1-BCART and ZIP2-BCART. To this end, we simulate a data set $\{(x_i, v_i, N_i)\}_{i=1}^n$ with $n = 5,000$ independent observations. Here $v_i \sim U(0, 1)$, $\mathbf{x}_i = (x_{i1}, x_{i2})$, with $x_{ik} \sim N(0, 1)$ for $k = 1, 2$. Moreover, $N_i \sim \text{ZIP}(p_i^{(\tau)}, \lambda(x_{i1}, x_{i2}) v_i)$, where

$$\lambda(x_1, x_2) = \begin{cases} 7 & \text{if } x_1 x_2 \leq 0, \\ 1 & \text{if } x_1 x_2 > 0, \end{cases}$$

and the probability of zero mass component is given as

$$p_i^{(\tau)} = \frac{\mu(x_{i1}, x_{i2})}{v_i^\tau + \mu(x_{i1}, x_{i2})}, \text{ with } \mu(x_{i1}, x_{i2}) \equiv 0.5,$$

and some $\tau \geq 0$ to be specified below. The data is split into two subsets, namely a training set with $n - m = 4,000$ observations and a test set with $m = 1,000$ observations.

In the above simulation setup, we include exposure in both the Poisson component and the zero mass component. In this way, it is not clear which of ZIP1-BCART and ZIP2-BCART will outperform the other. That being said, we could vary the value of τ to control the effect of exposure to the zero mass component. We shall consider two extreme cases, one with a very small τ and the other with a very large τ . More precisely, for a large τ we choose $\tau = 100$. In this case, since many v_i^τ will be small, we have that $p_i^{(\tau)}$ will be close to one, which implies that Poisson component should play a minor role in exposure modelling and thus we would expect that ZIP2-BCART has better ability to capture this. On the other hand, for a small value $\tau = 0.0001$, since many v_i^τ will be close to 1 we have that $p_i^{(\tau)}$ will be almost independent of v_i , which implies that zero mass component should play a minor role in exposure modelling and thus we would expect that ZIP1-BCART has better ability to capture this. We report DIC for these two cases in Table 11. The model performances on test data are listed in Table 12 for $\tau = 100$ and Table 13 for $\tau = 0.0001$. From these tables, we can confirm the above intuition that ZIP1-BCART should perform better for small τ and worse for large τ (compared to ZIP2-BCART). We conclude from this simulation study that the ZIP2-BCART works better in capturing the potential stronger effect of the exposure to the zero mass component, which is also illustrated in the real insurance data discussed below.

Table 11
DIC for ZIP-BCART with different values of τ on training data. Bold font indicates DIC selected model.

Model	DIC ($\tau = 100$)	DIC ($\tau = 0.0001$)
ZIP1-BCART (2)	3091	10515
ZIP1-BCART (3)	3055	10437
ZIP1-BCART (4)	2976	10273
ZIP1-BCART (5)	2997	10330
ZIP2-BCART (2)	2653	10924
ZIP2-BCART (3)	2637	10843
ZIP2-BCART (4)	2613	10685
ZIP2-BCART (5)	2627	10751

Table 12
Model performance on test data ($\tau = 100$) with bold entries determined by DIC (see Table 11).

Model	RSS(N)	RSS(N/ν) (in 10^{-5})	NLL	DS(N/ν)	Lift
ZIP1-BCART (2)	2423	3.06	1339	0.00281	1.01
ZIP1-BCART (3)	2417	2.98	1330	0.00259	1.33
ZIP1-BCART (4)	2376	2.20	1308	0.00209	1.78
ZIP1-BCART (5)	2333	2.63	1302	0.00219	1.81
ZIP2-BCART (2)	2072	2.76	1324	0.00234	1.06
ZIP2-BCART (3)	2069	2.57	1317	0.00207	1.46
ZIP2-BCART (4)	2056	2.02	1304	0.00179	1.97
ZIP2-BCART (5)	2049	2.06	1295	0.00189	2.08

Table 13
Model performance on test data ($\tau = 0.0001$) with bold entries determined by DIC (see Table 11).

Model	RSS(N)	RSS(N/ν)	NLL	DS(N/ν)	Lift
ZIP2-BCART (2)	6859	0.0093	4185	0.0080	1.02
ZIP2-BCART (3)	6648	0.0080	4092	0.0069	2.10
ZIP2-BCART (4)	6408	0.0060	3913	0.0050	3.40
ZIP2-BCART (5)	6320	0.0073	3853	0.0062	3.48
ZIP1-BCART (2)	6628	0.0079	3827	0.0072	1.07
ZIP1-BCART (3)	6535	0.0058	3763	0.0055	2.15
ZIP1-BCART (4)	6350	0.0027	3590	0.0024	3.45
ZIP1-BCART (5)	6282	0.0036	3543	0.0033	3.62

Table 14
Description of variables (*dataCar*).

Variable	Description	Type
numclaims	number of claims	numeric
exposure	in yearly units, between 0 and 1	numeric
veh_value	vehicle value, in \$10,000s	numeric
veh_age	vehicle age category, 1 (youngest), 2, 3, 4	numeric
agecat	driver age category, 1 (youngest), 2, 3, 4, 5, 6	numeric
veh_body	vehicle body, include 13 different types coded as HBACK, UTE, STNWG, HDTOP, PANVN, SEDAN, TRUCK, COUPE, MIBUS, MCARA, BUS, CONVT, RDSTR	character
gender	Female or Male	character
area	coded as A B C D E F	character

4.3. Real data analysis

We illustrate our methodology with a real insurance dataset, named *dataCar*, available from the library *insuranceData* in R; see Wolny-Dominiak and Trzesiok (2014) for details. This dataset is based on one-year vehicle insurance policies taken out in 2004 or 2005. There are 67,856 policies of which 93.19% made no claims. A summary of the variables used is given in Table 14. We split this dataset into training (80%) and test (20%) data sets, in doing so we keep the balance of zero and non-zero claims in both training and test data sets.

We shall apply the BCART models for claims frequency modelling introduced in Section 3 to training data, where we can use the three-step approach given in Table 1 to choose an optimal tree for each model (and also a global optimal one). We then assess the performance of these obtained trees on test data.

Running ANOVA-CART on the training data, we use cross-validation to select the tree size, which has 5 terminal nodes. We also run P-CART in the same way, again resulting in a tree with 5 terminal nodes, and this tree is shown in Fig. 4. Then, we apply P-BCART, NB1-BCART, NB2-BCART, ZIP1-BCART and ZIP2-BCART to the same data. Based on the knowledge learnt from CARTs above, we can tune the hyper-parameters γ, ρ , so that the algorithm will converge to a region of trees with number of terminal nodes around 5. Some of these, together with the effective number of

Table 15
Hyper-parameters, p_D (or q_D, r_D) and DIC on training data (*dataCar*). Bold font indicates DIC selected model.

Model	γ	ρ	p_D (or q_D, r_D)	DIC
P-BCART (4)	0.99	15	4.00	27948.8
P-BCART (5)	0.99	8	5.00	27943.8
P-BCART (6)	0.99	6	6.00	27944.4
NB1-BCART (4)	0.99	15	7.98	26002.4
NB1-BCART (5)	0.99	7	9.96	25892.0
NB1-BCART (6)	0.99	6	11.96	25945.2
NB2-BCART (4)	0.99	15	7.99	25925.7
NB2-BCART (5)	0.99	6	9.98	25846.4
NB2-BCART (6)	0.99	5	11.97	25885.6
ZIP1-BCART (4)	0.99	10	8.05	25688.4
ZIP1-BCART (5)	0.99	5	9.85	25674.1
ZIP1-BCART (6)	0.99	3	12.00	25678.3
ZIP2-BCART (4)	0.99	10	7.99	25654.3
ZIP2-BCART (5)	0.99	4	9.91	25632.5
ZIP2-BCART (6)	0.99	3	11.93	25641.4

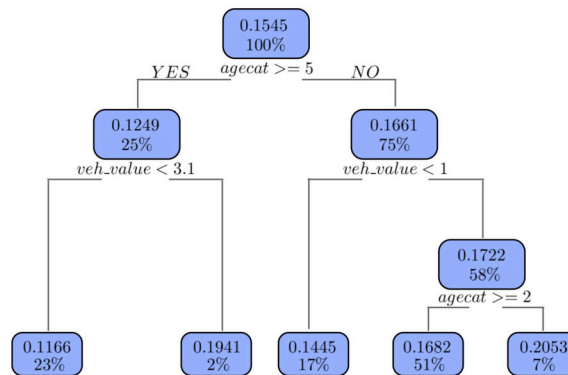


Fig. 4. Tree from P-CART. Numbers at each node give the estimated frequency and the percentage of observations.

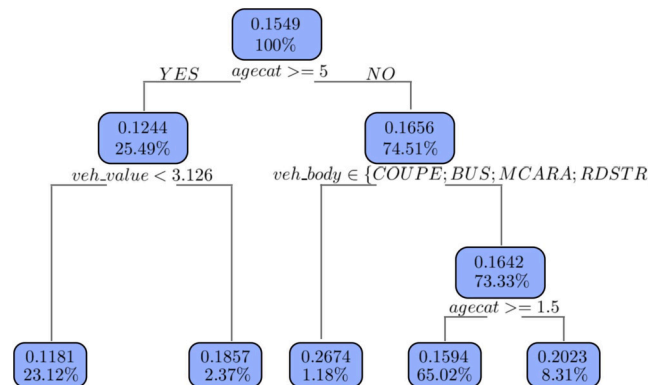


Fig. 5. Optimal tree from ZIP2-BCART. Numbers at each node give the estimated frequency and the percentage of observations.

parameters and DIC, are shown in Table 15. We see from this table that all the effective numbers of parameters are reasonable for the model used to fit the data. We conclude from the DIC that all of these BCART models select an optimal tree with 5 terminal nodes using the three-step approach, and among these the one from ZIP2-BCART, with the smallest DIC (= 25632.5), should be chosen as the global optimal tree to characterize the data.

It is interesting to check whether there are similarities in the trees obtained from different models, including the P-CART, particularly as they all have 5 terminal nodes. For the tree from P-CART illustrated in Fig. 4, the variable “agecat” is first used and then “veh_value”, followed by “agecat” again. The tree from P-BCART (not shown here) also uses “agecat” first, but in the following steps, it uses “veh_value” and “veh_body”. The trees from NB1-BCART and NB2-BCART look very similar, and both of them use “gender” first and then use “agecat”, “veh_value” and “veh_body”. Further, the trees from ZIP1-BCART and ZIP2-BCART have the same tree structure and select the same splitting variables as the tree from P-BCART, while the split values/categories are slightly different. The optimal tree from ZIP2-BCART is displayed in Fig. 5, where the estimated frequency (i.e., the first figure in each node) is calculated through (14) for the ZIP2 model with unit exposure. Comparing the two trees in Figs. 4 and 5 we see that ZIP2-BCART model can identify a more risky group (i.e., the one with estimated frequency equal to 0.2674). Moreover, for comparison we also use GLM to fit the data. We find that only the variables “agecat” and “veh_body” are significant, in which we also use the interactions between these two

Table 16
Model performance on test data (*dataCar*) with bold entries determined by DIC (see Table 15).

Model	RSS (N)	RSS (N/ν)	NLL	DS(N/ν)	Lift	Time (s)	Memory (MB)
GLM	1057.029	-	5532.37	-	-	1.15	115
ANOVA-CART (5)	1054.061	0.0205	5514.06	0.0700	1.83	2.05	98
P-CART (5)	1042.295	0.0185	5476.43	0.0681	1.97	2.13	98
P-BCART (4)	1042.221	0.0172	5473.90	0.0680	1.74	317.61	364
P-BCART (5)	1042.211	0.0167	5472.86	0.0602	2.26	291.28	378
P-BCART (6)	1042.205	0.0171	5472.27	0.0632	2.29	325.10	581
NB1-BCART (4)	1041.129	0.0168	5470.12	0.0445	1.80	413.95	628
NB1-BCART (5)	1041.109	0.0159	5469.00	0.0372	2.46	403.84	569
NB1-BCART (6)	1041.103	0.0162	5468.51	0.0413	2.57	459.70	689
NB2-BCART (4)	1041.127	0.0155	5470.01	0.0416	1.85	431.90	642
NB2-BCART (5)	1041.102	0.0144	5468.68	0.0352	2.50	441.82	661
NB2-BCART (6)	1041.094	0.0151	5468.35	0.0390	2.58	492.19	721
ZIP1-BCART (4)	1041.102	0.0150	5469.07	0.0383	1.91	548.29	827
ZIP1-BCART (5)	1041.087	0.0138	5468.39	0.0316	2.56	524.84	792
ZIP1-BCART (6)	1041.075	0.0142	5468.02	0.0362	2.60	569.21	889
ZIP2-BCART (4)	1041.054	0.0145	5468.25	0.0279	2.20	561.98	840
ZIP2-BCART (5)	1041.038	0.0136	5468.01	0.0241	2.72	570.40	851
ZIP2-BCART (6)	1041.025	0.0141	5467.81	0.0271	2.79	589.24	892

variables. In conclusion, though the variables used for different models can differ slightly, there seems to be a consensus that “*agecat*”, “*veh_value*” and “*veh_body*” are relatively significant variables and “*gender*”, “*veh_age*” and “*area*” are less significant.

Now, we apply the trees to the test data. The performances are given in Table 16. We also include the commonly used GLM, for which the performance looks not as good as the tree models. From the table, we can conclude that for each of the BCART models the tree with 5 terminal nodes that is selected by DIC performs better, in terms of $RSS(N/\nu)$ and $DS(N/\nu)$, than the trees with either smaller or larger number of terminal nodes. This confirms that the proposed three-step approach for the tree model selection in each type of models based on DIC works well in real data. Moreover, all the performance measures give the same ranking of models (from best to worst) as follows:

ZIP2-BCART, ZIP1-BCART, NB2-BCART, NB1-BCART, P-BCART, P-CART, ANOVA-CART, GLM.

This ranking is, to some extent, consistent with the conclusions from the simulation examples and as expected. We do not know the exact distribution of real insurance data, but we do know that it contains a high proportion of zeros, where the advantage of ZIP comes into play. Further, comparing NB and Poisson distributions, the former is able to handle over-dispersion, so their performance ranking is reasonable. Moreover, the ranking of two ZIP-BCART models and two NB-BCART models are also consistent with the conclusions of Lee (2020, 2021) where it is justified that the non-standard ways of dealing with exposures (i.e., ZIP2-BCART and NB2-BCART) should better fit real insurance data.

In addition to the performance measures, we also record the computation time (in seconds) and memory usage (in megabytes); see the last two columns of Table 16. All computations were performed on a laptop with Processor (3.5 GHz Dual-Core Intel Core i7) and Memory (16 GB 2133 MHz LPDDR3). Clearly, BCART models are far inferior to CARTs and GLM in these two respects and as the number of latent variables increases (from P-BCART to NB-CART to ZIP-BCART) these indicators become worse, but we think with such a large training data these are still acceptable and feasible to use in practice. We remark that there have been prior endeavours to address computing issues; see, e.g., Chipman et al. (2014); He et al. (2019); Sparapani et al. (2021). We believe these two indicators will be improved after our code is optimized in the future.

We conclude this section with some discussions on the *stability* of the proposed BCART models. Stability is a notion in computational learning theory of how the output of a machine learning algorithm is perturbed by small changes to its inputs. A stable learning algorithm is one for which the prediction does not change much when training data is modified slightly; see, e.g., Arsov et al. (2019) and references therein. CART models are known to be unstable. It is thus interesting to examine whether the proposed BCART models can be more stable. To this end, we propose the following approach to assess the stability of the P-CART and ZIP2-BCART (as the best) models.

- Randomly divide the data into two parts, 80% for training and 20% for testing.
- Randomly select 90% of training data for 20 times to construct 20 training subsets, named $Data_1, Data_2, \dots, Data_{20}$.
- Obtain the optimal tree from P-CART and ZIP2-CART, respectively, for each training set $Data_j, j = 1, \dots, 20$.
- Use the previously obtained trees to get predictions for test data. For each observation in test data, we will have 20 predictions from the 20 P-CART trees for which we calculate the variance, and do the same for the 20 ZIP2-BCART trees to get a variance.
- Calculate the mean (over the observations in test data) of those variances for P-CART and ZIP2-BCART, respectively.

Since variance can capture the amount of variability, we shall use the above obtained mean to assess the stability (in their predicting ability) of a tree-based model. Namely, the smaller the mean the more stable the model that was used to calculate it. We apply it to the *dataCar* insurance data, the calculated mean for P-CART is 9.319339×10^{-5} and for ZIP2-BCART is 6.896231×10^{-5} . This implies that ZIP2-BCART is more stable than P-CART. Additionally, we also compare the 20 trees from P-CART, where we can observe very different trees in terms of number of terminal nodes (ranging from 3 to 8) and splitting variables selected in the trees. Whereas, the 20 trees from ZIP2-BCART also show some stability in terms of number of terminal nodes (all around 5) and splitting variables selected. The same procedure has also been applied to other BCART models and the conclusions are almost the same. Therefore, we conclude from our studies that the proposed BCART models in this paper show some stability that the CART models may not possess.

5. Summary and discussions

This work proposes the use of BCART models for insurance pricing, and in particular, claims frequency prediction. These tree-based models can automatically perform variable selection and detect non-linear effects and possible interactions among explanatory variables. The obtained optimal trees are relatively accurate, stable and are straightforward to interpret by a visualization of the tree structure. These are desirable aspects for insurance pricing. We have introduced the framework of the BCART models and presented MCMC algorithms for general non-Gaussian distributed data where data augmentation may be needed in its implementation. We have included BCART models for Poisson, NB and ZIP distributions, which are the commonly used distributions for claims frequency. For the NB and ZIP models, we explored two different ways to deal with exposures. Remarkably, we conclude from the simulation examples and real data analysis that the non-standard ways of embedding exposures can provide us with better tree models, which is in line with the conclusions of Lee (2020, 2021). Furthermore, we introduced a tree model selection approach based on DIC, which has been seen to be an effective approach using both simulation examples and real insurance data. In particular, we conclude from the real insurance data analysis that the ZIP-BCART with exposure embedded in the zero mass component is the best candidate for claims frequency modelling. It is worth remarking that a general zero-inflated NB BCART can be implemented and may further improve the accuracy, but this will require more latent variables to be introduced and will make the convergence of MCMC algorithm harder/slower; see Murray (2021) for some insights.

Below we comment on potential further improvements of the BCART models for claims frequency modelling.

- In the MCMC algorithms we have only used four common proposals, namely, Grow, Prune, Change, and Swap, which have made the algorithm to quickly converge to a local optimal region. In order to make it better explore the tree space, other proposals such as those in Wu et al. (2007); Pratola (2016) can be suggested to improve the mixing of simulated trees. However, we suspect this will significantly increase the computational time, particularly, for high-dimensional large data set and for models requiring data augmentation. To mitigate this effect, we might consider to use a non-uniform choice of splitting variables in the tree prior so as to achieve a better variable selection, e.g., the Dirichlet prior proposed in Linero (2018).
- The proposed models have imposed several assumptions in order to simplify calculations. For example, we used conjugate prior for the terminal node distributions, and additional independence assumption as in (4) and (45). To further improve the analysis, it might be beneficial to incorporate different specifications of the prior for the same distribution scenario without using conjugate priors or independence, while this may require other techniques such as Laplace approximation (see Chipman et al. (2003)). We refer also to Chipman and McCulloch (2000) for an interesting incorporation of some hierarchical priors.
- We have proposed to use a single (optimal) tree induced from the BCART models for claims frequency prediction. The main reason for this choice is, as we discussed in the introduction, for ease of interpretation. Since stakeholders and regulators may not be statisticians who are able to understand very complex statistical models, a single tree offers intuitive and visual results to them. Although we have proposed an approach to find one single optimal tree, some sub-optimal trees (in the convergence region of the MCMC) which possess similar/different tree structures, may also be as informative as the single optimal tree and should not be simply ignored. Further research can be done in this direction to make better use of the posterior trees by clustering or merging them; see, e.g., Chipman et al. (2001); Banerjee et al. (2012).
- To further improve the accuracy of these Bayesian tree-based models we could explore BART for claims frequency modelling. The BART models are tree ensembles; each tree in BART only accounts for a small part of the overall fit, potentially improving the performance, but model interpretability needs to be explored before it can be used for insurance pricing. To this end, we believe some insights from Henckaerts et al. (2021) would be helpful.

In this paper, we have focused on insurance claims frequency. A natural next step is to construct a full insurance pricing BCART model, including both claims frequency and severity.

Declaration of competing interest

There is no competing interest.

Data availability

The data is in some R package.

Acknowledgement

We are thankful to the anonymous referee for their constructive suggestions which have led to a significant improvement of the manuscript.

References

- Arsov, N., Pavlovski, M., Kocarev, L., 2019. Stability of decision trees and logistic regression. Preprint. <https://arxiv.org/pdf/1903.00816.pdf>.
- Athreya, K.B., Ney, P.E., 2004. *Branching Processes*. Courier Corporation.
- Banerjee, M., Ding, Y., Noone, A.-M., 2012. Identifying representative trees from ensembles. *Statistics in Medicine* 31 (15), 1601–1616.
- Bleich, J., Kapelner, A., George, E.L., Jensen, S.T., 2014. Variable selection for BART: an application to gene regulation. *Annals of Applied Statistics* 8 (3), 1750–1781.
- Blier-Wong, C., Cossette, H., Lamontagne, L., Marceau, E., 2020. Machine learning in P&C insurance: a review for pricing and reserving. *Risks* 9 (1), 4.
- Breiman, L., Friedman, J., Stone, C.J., Olshen, R.A., 1984. *Classification and Regression Trees*. CRC Press.
- Bühlmann, H., Gisler, A., 2005. *A Course in Credibility Theory and Its Applications*, vol. 317. Springer.
- Celeux, G., Forbes, F., Robert, C.P., Titterton, D.M., 2006. Deviance information criteria for missing data models. *Bayesian Analysis* 1 (4), 651–673.
- Chipman, H., McCulloch, R.E., 2000. Hierarchical priors for Bayesian cart shrinkage. *Statistics and Computing* 10, 17–24.
- Chipman, H., George, E., McCulloch, R., 2003. Bayesian treed generalized linear models. *Bayesian Statistics* 7, 323–349.
- Chipman, H., George, E., Hahn, R., McCulloch, R., Pratola, M., Sparapani, R., 2014. Bayesian additive regression trees, computational approaches, Wiley StatsRef: Statistics Reference Online, pp. 1–23.

- Chipman, H.A., George, E.I., McCulloch, R.E., 1998. Bayesian CART model search. *Journal of the American Statistical Association* 93 (443), 935–948.
- Chipman, H.A., George, E.I., McCulloch, R.E., 2001. Managing multiple models. In: *International Workshop on Artificial Intelligence and Statistics*, PMLR, pp. 41–48.
- Chipman, H.A., George, E.I., McCulloch, R.E., 2002. Bayesian treed models. *Machine Learning* 48 (1), 299–320.
- Chipman, H.A., George, E.I., McCulloch, R.E., et al., 2010. BART: Bayesian additive regression trees. *Annals of Applied Statistics* 4 (1), 266–298.
- Denison, D.G., Mallick, B.K., Smith, A.F., 1998. A Bayesian CART algorithm. *Biometrika* 85 (2), 363–377.
- Denuit, M., Trufin, J., 2019. *Effective Statistical Learning Methods for Actuaries*. Springer.
- Denuit, M., Maréchal, X., Pitrebois, S., Walhin, J.-F., 2007. Actuarial Modelling of Claim Counts: Risk Classification, Credibility and Bonus-Malus Systems. John Wiley & Sons.
- Denuit, M., Charpentier, A., Trufin, J., 2021. Autocalibration and Tweedie-dominance for insurance pricing with machine learning. *Insurance: Mathematics and Economics* 101, 485–497.
- Diebolt, J., Robert, C.P., 1994. Estimation of finite mixture distributions through Bayesian sampling. *Journal of the Royal Statistical Society, Series B, Methodological* 56 (2), 363–375.
- Gelman, A., Hwang, J., Vehtari, A., 2014. Understanding predictive information criteria for Bayesian models. *Statistics and Computing* 24 (6), 997–1016.
- George, E.I., 1998. Bayesian model selection. In: *Encyclopedia of Statistical Sciences Update*, vol. 3.
- Green, P.J., 1995. Reversible jump MCMC computation and Bayesian model determination. *Biometrika* 82, 711–732.
- He, J., Yalov, S., Hahn, P.R., 2019. XBART: accelerated Bayesian additive regression trees. In: *The 22nd International Conference on Artificial Intelligence and Statistics*, PMLR, pp. 1130–1138.
- Henckaerts, R., Antonio, K., Clijsters, M., Verbelen, R., 2018. A data driven binning strategy for the construction of insurance tariff classes. *Scandinavian Actuarial Journal* 2018 (8), 681–705.
- Henckaerts, R., Côté, M.-P., Antonio, K., Verbelen, R., 2021. Boosting insights in insurance tariff plans with tree-based machine learning methods. *North American Actuarial Journal* 25 (2), 255–285.
- Hill, J., Linero, A., Murray, J., 2020. Bayesian additive regression trees: a review and look forward. *Annual Review of Statistics and Its Application* 7, 251–278.
- Hu, C., Quan, Z., Chong, W.F., 2022. Imbalanced learning for insurance using modified loss functions in tree-based models. *Insurance: Mathematics and Economics* 106, 13–32.
- Kapeller, A., Bleich, J., 2013. BartMachine: machine learning with Bayesian additive regression trees. arXiv preprint. arXiv:1312.2171.
- Kindo, B.P., Wang, H., Peña, E.A., 2016. Multinomial probit Bayesian additive regression trees. *Stat* 5 (1), 119–131.
- Lee, S.C., 2020. Delta boosting implementation of negative binomial regression in actuarial pricing. *Risks* 8 (1), 19.
- Lee, S.C., 2021. Addressing imbalanced insurance data through zero-inflated Poisson regression with boosting. *ASTIN Bulletin: The Journal of the IAA* 51 (1), 27–55.
- Lindholm, M., Lindskog, F., Palmquist, J., 2022. Local bias adjustment, duration-weighted probabilities, and automatic construction of tariff cells. In: *Duration-Weighted Probabilities, and Automatic Construction of Tariff Cells*. October 24, 2022.
- Linero, A.R., 2017. A review of tree-based Bayesian methods. *Communications for Statistical Applications and Methods* 24 (6), 543–559.
- Linero, A.R., 2018. Bayesian regression trees for high-dimensional prediction and variable selection. *Journal of the American Statistical Association* 113 (522), 626–636.
- Linero, A.R., Yang, Y., 2018. Bayesian regression tree ensembles that adapt to smoothness and sparsity. *Journal of the Royal Statistical Society, Series B, Statistical Methodology* 80 (5), 1087–1110.
- Linero, A.R., Sinha, D., Lipsitz, S.R., 2020. Semiparametric mixed-scale models using shared Bayesian forests. *Biometrics* 76 (1), 131–144.
- Liu, Y., Ročková, V., Wang, Y., 2021. Variable selection with ABC Bayesian forests. *Journal of the Royal Statistical Society, Series B, Statistical Methodology* 83 (3), 453–481.
- Meng, S., Gao, Y., Huang, Y., 2022. Actuarial intelligence in auto insurance: claim frequency modeling with driving behavior features and improved boosted trees. *Insurance: Mathematics and Economics* 106, 115–127.
- Meng, X.-L., Van Dyk, D.A., 1999. Seeking efficient data augmentation schemes via conditional and marginal augmentation. *Biometrika* 86 (2), 301–320.
- Murray, J.S., 2021. Log-linear Bayesian additive regression trees for multinomial logistic and count regression models. *Journal of the American Statistical Association* 116 (534), 756–769.
- Naya, H., Urioste, J.I., Chang, Y.-M., Rodrigues-Motta, M., Kremer, R., Gianola, D., 2008. A comparison between Poisson and zero-inflated Poisson regression models with an application to number of black spots in Corriedale sheep. *Genetics Selection Evolution* 40, 1–16.
- Nelder, J.A., Wedderburn, R.W., 1972. Generalized linear models. *Journal of the Royal Statistical Society. Series A. General* 135 (3), 370–384.
- Ohlsson, E., Johansson, B., 2010. *Non-life Insurance Pricing with Generalized Linear Models*, vol. 174. Springer.
- Prado, E.B., Parnell, A.C., Murphy, K., McJames, N., O’Shea, A., Moral, R.A., 2021. Accounting for shared covariates in semi-parametric Bayesian additive regression trees. arXiv preprint. arXiv:2108.07636.
- Pratola, M.T., 2016. Efficient Metropolis–Hastings proposal mechanisms for Bayesian regression tree models. *Bayesian Analysis* 11 (3), 885–911.
- Quan, Z., 2019. *Insurance Analytics with Tree-Based Models*. PhD thesis. University of Connecticut.
- Ročková, V., Van der Pas, S., et al., 2020. Posterior concentration for Bayesian regression trees and forests. *The Annals of Statistics* 48 (4), 2108–2131.
- Rodrigues, J., 2003. Bayesian analysis of zero-inflated distributions. *Communications in Statistics - Theory and Methods* 32 (2), 281–289.
- Saha, E., 2023. Theory of posterior concentration for generalized Bayesian additive regression trees. arXiv preprint. arXiv:2304.12505.
- Sparapani, R., Spanbauer, C., McCulloch, R., 2021. Nonparametric machine learning and efficient computation with Bayesian additive regression trees: the BART R package. *Journal of Statistical Software* 97, 1–66.
- Spiegelhalter, D.J., Best, N.G., Carlin, B.P., Van Der Linde, A., 2002. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society, Series B, Statistical Methodology* 64 (4), 583–639.
- Spiegelhalter, D.J., Best, N.G., Carlin, B.P., Van der Linde, A., 2014. The deviance information criterion: 12 years on. *Journal of the Royal Statistical Society, Series B, Statistical Methodology* 76 (3), 485–493.
- Tanner, M.A., Wong, W.H., 1987. The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* 82 (398), 528–540.
- Therneau, T., Atkinson, B., 2019. rpart: recursive partitioning and regression trees, R package version 4.1-15. [Online]. Available. <https://CRAN.R-project.org/package=rpart>.
- Van Dyk, D.A., Meng, X.-L., 2001. The art of data augmentation. *Journal of Computational and Graphical Statistics* 10 (1), 1–50.
- Watanabe, S., Opper, M., 2010. Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research* 11 (12).
- Wolny-Dominiak, A., Trzeciok, M., 2014. Insurancedata: a collection of insurance datasets useful in risk classification in non-life insurance, R package version 1.0. [Online]. Available. <https://CRAN.R-project.org/package=insuranceData>.
- Wu, Y., Tjelmeland, H., West, M., 2007. Bayesian CART: prior specification and posterior simulation. *Journal of Computational and Graphical Statistics* 16 (1), 44–66.
- Wuthrich, M., Merz, M., 2022. *Statistical Foundations of Actuarial Learning and Its Applications*. Springer Actuarial. Open Access. <https://link.springer.com/book/10.1007/978-3-031-12409-9>. Available at SSRN. <https://ssrn.com/abstract=3822407>.
- Wüthrich, M.V., 2020. Bias regularization in neural network models for general insurance pricing. *European Actuarial Journal* 10 (1), 179–202.
- Wüthrich, M.V., 2022. *Non-life Insurance: Mathematics & Statistics*. Available at SSRN 2319328.
- Wüthrich, M.V., 2022. The balance property in neural network modelling. *Statistical Theory and Related Fields* 6 (1), 1–9.
- Wüthrich, M.V., Buser, C., 2023. Data analytics for non-life insurance pricing. January 9, 2023. Available at SSRN. <https://ssrn.com/abstract=2870308>.
- Zhou, M., Li, L., Dunson, D., Carin, L., 2012. Lognormal and gamma mixed negative binomial regression. In: *Proceedings of the International Conference on Machine Learning*. International Conference on Machine Learning. In: NIH Public Access, vol. 2012, p. 1343.